



TESI DI DOTTORATO DI RICERCA IN CO-TUTELA

UNIVERSITÀ DEGLI STUDI DI MILANO

UNIVERSITÉ CATHOLIQUE DE LOUVAIN

Programma di Dottorato in Fisica, Astrofisica e
Fisica Applicata, Ciclo XXXVIII

Degree of Doctor in Sciences

Precision tests of the Standard Model: the Drell-Yan processes

Settore Scientifico Disciplinare FIS/02

Tommaso ARMADILLO

Matr. n. R14738

ORCID n. 0000-0001-8609-1334

Supervisors: Prof. Fabio MALTONI, Prof. Alessandro VICINI

PhD Program Coordinator: Prof. Aniello Mennella

Academic Year 2025/2026

Commission of the final examination:

Supervisors

Prof. Fabio MALTONI, Université catholique de Louvain,
Prof. Alessandro VICINI, Università degli studi di Milano,

External Members:

Prof. Jonas M. Lindert, University of Sussex,
Prof. Stefano Pozzorini, Universität Zürich,

Internal Members:

Prof. Philippe Ruelle, Université catholique de Louvain,
Prof. Gauthier Durieux, Université catholique de Louvain.

Private Defense:

Date: 10/07/2026

Place: Université catholique de Louvain, Centre for Cosmology and Particle Physics Phenomenology,
Louvain-la-Neuve, Belgium.

Public Defense:

Date: 03/09/2026

Place: Università degli Studi di Milano, Dipartimento di Fisica, Milano, Italy.

Abstract

The Standard Model (SM) stands as one of the most successful and extensively tested theoretical framework for describing fundamental particles and their interactions. To date, experimental measurements across a vast range of energy scales have shown outstanding agreement with theoretical predictions. Despite these experimental triumphs, the theory remains fundamentally incomplete. For example, it fails to account for established physical realities such as dark matter, the matter-antimatter asymmetry of the universe, neutrino masses, and a quantum description of gravity. New physics typically manifests itself in experimental data as resonances in differential distributions. However, the absence of such direct observations indicates that, if present, it resides at energy scales beyond the kinematic reach of current colliders. While the on-shell peaks of these resonances may not be experimentally accessible, their low-energy tails can still manifest as subtle deviations in the TeV region. For this reason, the frontier of high-energy physics has shifted towards precision tests. During the High-Luminosity phase of the LHC, these high-energy regions are projected to be measured with percent or sub-percent level accuracy. To ensure a statistically meaningful comparison between experimental data and theoretical predictions, the uncertainties on the latter must be systematically reduced to the same level.

This thesis contributes to the advancement of these theoretical predictions by calculating higher-order perturbative corrections, with a specific emphasis on the electroweak (EW) sector of the SM. To this end, we developed OCEANN, a novel computational framework that lays down the pipeline from the generation of Feynman diagrams to the numerical evaluation of ultraviolet-renormalized and infrared-subtracted two-loop amplitudes. Within OCEANN, we introduced ABISS, a MATHEMATICA package designed to compute the interference between two-loop and tree-level amplitudes. While to overcome the severe computational bottlenecks associated with evaluating complex multi-scale Feynman integrals, we developed SEASYDE, a MATHEMATICA package that implements the series expansion approach. Crucially, this newly developed computational tools natively support complex masses, allowing for the stable and precise evaluation of amplitudes in resonant processes.

We applied this framework to compute the mixed QCD-EW two-loop amplitudes for both Neutral Current (NC) and Charged Current (CC) Drell-Yan (DY) processes. Alongside this, we calculated the complete two-loop QED corrections to NC-DY, an effort motivated by the need to deeply understand the infrared structure of the full two-loop EW amplitude. These calculations represents a theoretical breakthrough for the LHC precision program, effectively opening the door towards complete two-loop EW computations.

Subsequently, we implemented the two-loop mixed amplitude for NC-DY into the Monte-Carlo integrator MATRIX and we conducted a detailed phenomenological study. This allowed us to provide precise differential cross-sections across both the resonant and high invariant-mass regions. We also performed a comprehensive comparison of our findings against an independent calculation available in literature.

Finally, we bridge the gap between precision phenomenology and Beyond the Standard Model physics searches by interpreting potential small deviations within the framework of the Standard Model Effective Field Theory (SMEFT). Utilizing the SMEFT framework, we integrate precise theoretical inputs into global SMEFT analyses to rigorously constrain the current parameter space of new physics. Furthermore, we extend this analytical framework to evaluate the sensitivity of future high-energy colliders. By projecting the foreseen increase in experimental precision at upcoming facilities into our global fits, we systematically quantify how future measurements will enhance our ability to isolate indirect signatures of the presence of new physics.

Associated Publications:

- [1] Tommaso Armadillo, Roberto Bonciani, Simone Devoto, Narayan Rana, Alessandro Vicini, "Two-loop mixed QCD-EW corrections to neutral current Drell-Yan", arXiv:2201.01754, JHEP, vol. 5, pp. 072, 2022.
- [2] Tommaso Armadillo, Roberto Bonciani, Simone Devoto, Narayan Rana, Alessandro Vicini, "Evaluation of Feynman integrals with arbitrary complex masses via series expansions", arXiv:2205.03345, Comput. Phys. Commun., vol. 282, pp. 108545, 2023.
- [3] Tommaso Armadillo, Roberto Bonciani, Simone Devoto, Narayan Rana, Alessandro Vicini, "Two-loop mixed QCD-EW corrections to charged current Drell-Yan", arXiv:2405.00612, JHEP, vol. 07, pp. 265, 2024.
- [4] Tommaso Armadillo, Roberto Bonciani, Luca Buonocore, Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, Narayan Rana, Alessandro Vicini, "Mixed QCD-EW corrections to the neutral-current Drell-Yan process", arXiv:2412.16095, JHEP, vol. 07, pp. 141, 2025.
- [5] Tommaso Armadillo, "Evaluating Feynman Integrals through differential equations and series expansions", arXiv:2502.14742, Eur.Phys.J.ST 234 (2026) 26, 7989-8004, 2025.
- [6] Tommaso Armadillo, Simone Devoto, Michele Dradi, Alessandro Vicini, "Towards the two-loop electroweak corrections to the Drell-Yan process: the infrared structure", arXiv:2511.20365.
- [7] Tommaso Armadillo, Eugenia Celada, Jaco ter Hoeve, Fabio Maltoni, Luca Mantani, Juan Rojo, Alejo N. Rossia, Simone Tentori, Marion O.A. Thomas, Eleni Vryonidou, "New physics reach through precision at future colliders: a multi-pronged approach", arXiv:2604.16596.
- [8] Tommaso Armadillo, Simone Devoto, Michele Dradi, Alessandro Vicini, "Towards the two-loop electroweak corrections to the Drell-Yan process: the complete fermionic contributions", arXiv:2605.28667.

The series expansion approach and our original package SEASYDE for the evaluation of Feynman integrals, presented in [2, 5] will be discussed in Chapter 3. In Chapter 5 we will focus on the evaluation of two-loop mixed QCD-EW corrections to neutral- [1] and charged-current [3] Drell-Yan and the two-loop QED amplitude for neutral-current Drell-Yan [6]. In Chapter 6 we will present phenomenological results for neutral-current Drell-Yan [4] at NNLO QCD-EW level of accuracy. Finally, in Chapter 7 we present our global fits in the context of Standard Model Effective Field Theory [7].

Contents

Introduction	1
1 The Standard Model of particle physics	5
1.1 Fundamental constituents of nature	6
1.2 Lagrangian of the Standard Model	6
1.2.1 Lagrangian of the electroweak SM in the symmetric basis	6
1.2.2 Lagrangian of the electroweak SM in the physical basis	8
1.2.3 Quantization	10
1.2.4 Inclusion of strong interaction	11
1.3 Background-field method	12
1.4 Ward identities	13
1.5 Problems of the Standard Model	14
1.6 Standard Model Effective Field Theory	16
2 Amplitudes	19
2.1 UV renormalization	21
2.1.1 Renormalization of parameters and fields	21
2.1.2 Renormalization conditions	22
2.1.3 Example diagrams	23
2.1.4 Counter-term expressions	24
2.1.5 Electric charge renormalization	27
2.1.6 Tadpole renormalization	28
2.2 Input schemes	29
2.2.1 $\alpha(0)$ scheme	30
2.2.2 G_μ scheme	30
2.3 Complex Mass Scheme	31
2.3.1 Impact on renormalization	31
2.3.2 Impact on input schemes	32
2.4 γ^5 problem	32
2.4.1 HVBM's prescription	33
2.4.2 Kreimer's prescription	34
2.4.3 d -dimensional Feynman integrals with four-dimensional scalar numerators	35
3 Feynman integrals	41
3.1 Definitions	41
3.1.1 Example: two-loop massless triangle	42
3.2 Integration-by-parts identities	43
3.3 Differential equations	44
3.3.1 The method	44

3.3.2	Example: 1L QED vertex	45
3.3.3	Boundary conditions	46
3.3.4	Epsilon expansion	46
3.3.5	Canonical basis	47
3.4	Series expansion approach	48
3.4.1	Bottom-to-top approach and solution of equations	48
3.4.2	Frobenius method	49
3.4.3	Variation of parameters	49
3.4.4	Example: 1L QED Vertex	50
3.4.5	Coupled equations	51
3.4.6	Variation of parameters for systems	53
3.5	Analytic continuation	53
3.6	Mass expansion	55
3.7	SeaSyde	57
3.7.1	Comparison with DIFFEXP	58
3.7.2	Checks and performance	59
4	From amplitudes to physical observables	63
4.1	Factorization theorem	63
4.2	NLO calculations	65
4.3	NNLO calculations	67
4.3.1	q_T -subtraction	67
4.4	Automatization of NNLO calculations	69
4.4.1	Two-loop amplitudes	69
4.4.2	Phase-space integration	71
5	Two-loop amplitudes for Drell-Yan	73
5.1	Two-loop mixed corrections to NC-DY	76
5.1.1	Framework of the calculation	76
5.1.2	Ultraviolet renormalisation	78
5.1.3	Generation of the full amplitude and evaluation of the interference terms	80
5.1.4	Integral families and the representation of the result in terms of MIs	82
5.1.5	Infrared singularities and universal pole structure	83
5.1.6	The UV-renormalised IR-subtracted hard function	85
5.1.7	Semi-analytical solution of MIs via series expansion	86
5.1.8	Numerical results	88
5.2	Two-loop mixed corrections to CC-DY	89
5.2.1	Generation of the full amplitude and evaluation of the interference terms	89
5.2.2	Integral families and the representation of the result in terms of MIs	91
5.2.3	Numerical evaluation of MIs via series expansions	94
5.2.4	Additional comments on the numerical evaluation of MIs	95
5.3	Two-loop EW corrections to NC-DY: IR structure	96
5.3.1	Infrared singularities and universal pole structure	96
5.3.2	The process	98
5.3.3	Ultraviolet renormalization	100
5.3.4	Explicit QED results and control over the cancellation of the IR poles	101

5.3.5	Impact of a change of scheme	103
5.4	Full Two-loop EW corrections: prospects	105
5.4.1	Diagrams	105
5.4.2	Integrals and differential equations	106
6	Phenomenology	109
6.1	Computational details	109
6.2	Phenomenological results	110
6.2.1	Resonant region	111
6.2.2	High invariant-mass region	114
6.2.3	Forward–Backward asymmetry	115
6.3	Comparison with massless calculation	119
7	Constraining new physics with SMEFT global fits	123
7.1	The framework of SMEFT global fits	123
7.1.1	SMEFT code	124
7.1.2	SMEFT Operator Basis	125
7.1.3	Experimental datasets and Drell-Yan	127
7.1.4	Projections for lepton colliders	131
7.1.5	Theoretical uncertainties	134
7.2	Global SMEFT analysis	134
7.2.1	Baseline results	135
7.2.2	Impact of quadratic EFT corrections	143
7.2.3	Impact of theoretical uncertainties	148
	Conclusion	151
	Acknowledgements	153
A	Renormalisation relations	157
A.1	Electric charge renormalisation within the $\overline{MS}$ scheme	157
A.2	Fermion mass and wave-function on-shell renormalisation	158
B	Photon vacuum polarization	159
	Bibliography	161

Introduction

“It is to know myself better and to find my own dimension that I have climbed impossible mountains. I did it driven by the beauty of alpine nature, by the challenge, and by the pleasure of knowing.”

Walter Bonatti, *Italian alpinist*.

The Standard Model of particle physics represents one of the most successful and extensively tested theoretical framework for describing fundamental particles and their interactions. Formulated over the second half of the twentieth century, it provides a remarkably elegant and predictive framework for describing the fundamental constituents of matter and three of the four known fundamental forces, namely electromagnetism, the weak and the strong nuclear force. For decades, the Standard Model has withstood intense experimental testing. From the precise measurements of electroweak (EW) observables at LEP [9] to the historic discovery of the Higgs boson at the Large Hadron Collider in 2012 [10, 11] whose existence has been predicted in 1964 by Peter Higgs and, independently, by Robert Brout and François Englert [12–15].

However, despite its great experimental successes, the Standard Model is fundamentally incomplete. For example, it fails to provide a quantum description of gravity. It does not offer any viable candidate for either dark matter or dark energy, which constitute up to 95% of the universe. It cannot explain the observed baryon asymmetry of the universe, nor the non-zero masses of neutrinos. These missing pieces therefore suggests that the Standard Model is not the ultimate theory of nature, but rather a low-energy effective approximation of a more fundamental theory. The primary objective of the high-energy physics community, is to uncover this Beyond the Standard Model physics.

In the early days of the Large Hadron Collider physics program, there was a widespread optimism for the direct discovery of new heavy resonances, such as supersymmetric particles or heavy gauge bosons. However, as the collider pushed to higher and higher center-of-mass energies and accumulated unprecedented amount of data, no such direct signals have emerged. The absence of new resonant phenomena at the TeV scale has induced a profound paradigm shift in the field: the transition from an era of discovery to an era of precision.

If new physics lies at energy scales significantly higher than those directly accessible by current colliders, its effects will show up as extremely small deviations in the kinematic distributions of known particles. To confidently claim a discovery based on these deviations, the theoretical uncertainties of our Standard Model predictions must be strictly smaller than the experimental ones. However, since the High-Luminosity phase at the Large Hadron Collider promises to reduce statistical experimental errors to the sub-percent level for many key observables [16, 17], the burden of discovery now rests heavily on the theoretical community. Achieving sub-percent theoretical accuracy on differential distributions requires extending perturbative calculations to higher orders in the strong and electroweak coupling constants. At Leading Order (LO), theoretical predictions typically suffer from large theoretical uncertainties. By computing Next-to-Leading Order (NLO) and Next-to-Next-to-Leading (NNLO) Order corrections, we systematically reduce them, stabilizing the convergence of the perturbative series.

Numerically, the fine-structure constant α is of the same order of magnitude as the square of the strong coupling α_s at the EW scale ($Q \sim m_Z$). Consequently, NLO EW corrections, formally $\mathcal{O}(\alpha)$ with respect to the LO, are comparable in size to NNLO QCD corrections, formally $\mathcal{O}(\alpha_s^2)$. However, EW corrections exhibit a unique behavior in the high-energy regime [18, 19]. In the tails of kinematic distributions, such as the high invariant mass regions, the exchange of virtual massive gauge bosons generates large logarithmic enhancements known as EW Sudakov logarithms. These corrections can impact the cross-section by tens of percent at the TeV scale. Because these high-energy tails are precisely the regions where new physics is most likely to show up, failing to include higher order EW corrections can artificially hide the signals of new physics. Achieving sub percent precision on differential distributions therefore demands not only the dominant higher order QCD contributions, but also the careful combination with higher-order mixed QCD-EW and purely EW effects.

Progress in the last 15 years has allowed the perturbative calculation of tree-level and one-loop amplitudes, together with their combination for predictions at NLO, to be completely automated both for QCD and EW loops. Nowadays, there exist multiple frameworks that perform all the calculations in a completely automatic way, e.g. MadGraph5_aMC@NLO [20]. Methods and techniques to compute NNLO corrections, on the other hand, are still under intense development [21]. Many results exist, yet obtained through dedicated and not-yet systematized approaches. In particular, the calculation of NNLO corrections, for a general scattering process, requires the inclusion of the contributions from different processes, namely double-real, real-virtual and double-virtual, which need first to be evaluated and then combined into infrared-safe observable. The evaluation of the double-virtual correction, however, is far from being automatized and a general standard workflow is not yet established.

The calculation of these higher-order corrections is a formidable mathematical and computational challenge. The workflow of a precision calculation begins with the algorithmic generation of thousands of Feynman diagrams. While this step is fully automated, the subsequent evaluation of the resulting multi-loop Feynman integrals constitutes the primary bottleneck of modern calculations [22]. Multi-loop amplitudes, indeed, diverge in exactly four dimensions due to ultraviolet and infrared singularities. To regularize these divergences, the integrals are evaluated in $d = 4 - 2\epsilon$ dimensions using dimensional regularization. The total number of d -dimensional scalar integrals grows exponentially with the perturbative order, reaching up to hundreds of thousands of integrals for state-of-the-art calculations. To render the problem tractable, we employ sophisticated algebraic techniques, leveraging Integration-by-Parts identities to reduce the vast system of integrals into a minimal, linearly independent basis known as Master Integrals. Evaluating these Master Integrals requires cutting-edge mathematical machinery. The modern standard relies on formulating and solving systems of differential equations with respect to the kinematic invariants of the scattering process [23]. Solutions are frequently expressed in terms of iterated integrals, such as multiple polylogarithms, or, in the presence of massive internal loops, more complex structures like elliptic integrals may arise. Mastering and automating these techniques is one of the central themes of this thesis.

The journey from a theoretical prediction to a physical constraint on new physics follows a standard pipeline, which forms the structural backbone of this thesis. In Chapter 1, we review the foundational aspects of the Standard Model of particle physics. We detail its theoretical formulation, introduce the background-field method for its quantization, and highlight some of the problems that motivate the search for new physics. Finally, we introduce the Standard Model Effective Field Theory as a framework to parametrize these new physics effects in a model-independent way. In Chapter 2, we introduce the core mathematical objects required for computing higher-order corrections: scattering amplitudes. We provide their formal definition and address the primary technical challenges associated with their

computation. Beginning with ultraviolet renormalization, we review its basic concepts and extend its formulation to the two-loop mixed QCD-EW and two-loop EW levels of accuracy. This extension, which was previously only partially available in the literature, has been presented in detail in [8]. We then review EW input schemes, the complex mass scheme, which is essential for accurately describing intermediate unstable particles, and the intricacies of treating γ^5 in d dimensions. In Chapter 3, we tackle the numerical evaluation of Feynman integrals. After providing the basic definitions and reviewing the method of differential equations, we focus extensively on the series expansion approach [5]. This powerful technique enables the computation of Feynman integrals even when analytical expressions in terms of known classes of functions are entirely unavailable. Finally, we present the Mathematica package SEASYDE [2]. Thanks to an original algorithm for performing the analytic continuation of results in the complex plane, this package is capable of evaluating Feynman integrals with arbitrary complex masses, thereby allowing us to consistently perform calculations within the complex mass scheme. In Chapter 4, we link the computed amplitudes to physical observables. Using the factorization theorem, we isolate the perturbative hard-scattering process from the non-perturbative structure of the proton. By employing infrared subtraction schemes and Monte Carlo integration techniques, we convert the finite matrix elements into fully differential fiducial cross-sections at the NLO and NNLO levels of accuracy. We then present OCEANN, our novel computational framework designed to automate the calculation of two-loop amplitudes. In Chapter 5, we focus on the Drell-Yan (DY) process, whose exceptionally clean leptonic final state provides the ultimate testing ground for precision physics. In particular, we present our calculations of the mixed QCD-EW two-loop amplitudes for Neutral-Current (NC) DY [1] and Charged-Current (CC) DY [3], as well as the complete two-loop QED corrections to NC-DY [6]. We conclude the chapter by discussing the prospects for achieving a full two-loop EW calculation. In Chapter 6, we present the implementation of our mixed QCD-EW amplitude for NC-DY into the Monte-Carlo integrator MATRIX, enabling phenomenological studies at the mixed NNLO QCD-EW level of accuracy [4]. We provide results in both the resonant and high-invariant-mass regions for various observables, and we present a detailed comparison with an independent calculation available in the literature. Finally, in Chapter 7, we utilize the SMEFT framework to integrate precise theoretical inputs into global SMEFT analyses, rigorously constraining the current parameter space of new physics. Furthermore, we extend this analytical framework to evaluate the sensitivity of future high-energy colliders [7]. By projecting the anticipated increase in experimental precision at upcoming facilities into our global fits, we systematically quantify how future measurements will enhance our ability to isolate indirect signatures of new physics.

In summary, this thesis presents a comprehensive journey through the precision frontier. By advancing the computation of multi-loop amplitudes and refining the phenomenological predictions for the Drell-Yan process, we improve the theoretical tools required to interrogate the Standard Model at unprecedented levels of accuracy, laying the foundations for the discovery of the fundamental laws of nature.

The Standard Model of particle physics

“Why do you want to climb Mount Everest?”

“Because it’s there.”

George Mallory, *English alpinist*.

Ever since humans began to think critically, we have been driven by a deep curiosity to understand the nature of reality. At the heart of this quest is a timeless question: *What is matter made of?* For millennia, philosophers and scientists have tried to identify the basic building blocks that constitute the physical world, striving to reduce the complexity of the universe to a set of elementary components.

The earliest recorded attempt to provide a systematic answer to this question emerged in Ancient Greece with the philosopher Democritus. In the 5th century BCE, he formulated the theory of atomism, proposing that the universe was composed of small, indivisible, and indestructible particles known as atoms¹, moving through an infinite void. For Democritus, these atoms were the fundamental constituents of all matter, varying only in shape, size, and arrangement. While his ideas were purely philosophical and lacked empirical verification, they represented a profound conceptual leap: the realization that the macroscopic world is supported by a microscopic substructure.

The scientific progress, however, has long since surpassed the atomism of antiquity. Through centuries of experiments and theoretical development, from the discovery of the electron to the formulation of quantum mechanics, our understanding of the subatomic world has deepened exponentially. We now know that the “uncuttable” atoms of Democritus are themselves composite structures, divisible into nuclei and electrons, and that the nuclei are further composed of protons and neutrons.

Today, the most robust and precise theoretical framework describing these fundamental constituents is the Standard Model (SM) of particle physics. It represents the culmination of modern physics, describing the universe not as a collection of geometric shapes in a void, but as a complex interplay of quantum fields. The SM successfully classifies all known elementary particles, namely quarks and leptons, and describes three of the four fundamental forces governing their interactions: the electromagnetic, weak, and strong nuclear forces. As the current apex of our understanding, the SM stands as one of the best theory humanity has produced to answer that primitive question, providing a mathematically consistent set of predictions that has withstood decades of rigorous experimental testing. This Chapter aims to provide a review of the theoretical foundations of the SM of particle physics. The discussion presented here is intended to establish the notation and introduce the fundamental concepts that will be used throughout the remainder of this thesis. For a more comprehensive treatment of the SM, quantum field theory, and gauge symmetries, we refer the reader to the extensive literature and classic textbooks available on the subject, such as [24–27].

¹From the ancient Greek *átomos*, not divisible.

1.1 Fundamental constituents of nature

In the framework of Quantum Field Theories (QFT), a fundamental object is defined as a point-like particle with no internal substructure, and, unlike the atoms or the nuclei, cannot be divided into smaller components. All visible matter in the universe is constructed from such fundamental particles known as fermions, which are characterized by having spin- $\frac{1}{2}$. These fermions are categorized into two distinct categories based on how they interact: quarks and leptons. Both quarks and leptons are further organized into three distinct generations (or families), where each successive generation consists of increasingly massive particles that otherwise share identical quantum numbers and interactions. The primary distinction between these two groups lies in their susceptibility to the strong nuclear force. Quarks possess a property known as color charge, making them subject to the strong interaction, whereas leptons do not. This property leads to a phenomenon called confinement: quarks are never found in isolation. Instead, they bind together to form composite particles called hadrons². Hadrons are then further classified into baryons, which are composed of three quarks (such as the proton and neutron), and mesons, which are unstable particles made of a quark and an anti-quark pair.

Interactions between fermions are built upon the fundamental principle of gauge symmetries. The SM is a non-Abelian gauge theory, invariant under the group $G = SU(3)_c \times SU(2)_L \times U(1)_Y$. This mathematical structure mandates the existence of force-carrying particles, known as gauge bosons. These spin-1 particles appear naturally from the theory, acting as the mediators for three fundamental forces (gravity being the notable exception). The color symmetry $SU(3)_c$ governs the strong interaction, mediated by a set of eight massless gluons. Meanwhile, the electroweak (EW) sector combines the $SU(2)_L$ and $U(1)_Y$ symmetries, giving rise to the photon (γ), which carries the electromagnetic force, and the $W^\pm$ and Z bosons, which mediate the weak nuclear force.

A critical distinction exists between these mediators. While the photon and gluons are massless, the W and Z bosons possess a significant mass, restricting the weak force to extremely short distances. This presents a theoretical challenge: strict gauge invariance normally forbids mass terms for these bosons, as well as for the fermions, contradicting experimental measurements. The resolution to this conflict can be found in the Brout-Englert-Higgs mechanism [12–15]. The inclusion of the Higgs field spontaneously breaks the electroweak symmetry, providing the W and Z bosons and the fermions with a mass, while leaving the photon massless. The associated particle, the spin-0 Higgs boson, was the final piece of the Standard Model puzzle to be discovered. The specific details of this symmetry breaking and mass generation will be explored in the subsequent section.

1.2 Lagrangian of the Standard Model

The SM is a gauge QFT which is symmetric under the group $G = SU(3)_c \times SU(2)_L \times U(1)_Y$. In particular, it unifies the strong interaction, described by QCD [28–31], with the electromagnetic and weak interactions, described by the Glashow-Salam-Weinberg model [32–34]. We start formulating the theory in a way in which the gauge symmetry is manifest and then, following a spontaneous symmetry breaking mechanism, we will switch to a representation in terms of physical fields.

1.2.1 Lagrangian of the electroweak SM in the symmetric basis

The electroweak sector of the Standard Model is defined by the gauge symmetry group $G_{EW} = SU(2)_L \times U(1)_Y$, representing the unification of the weak and electromagnetic interactions. The coupling strengths

²From the ancient Greek *hadrós*, large, massive.

associated with these groups are denoted by g_2 for the weak isospin $SU(2)_L$ and g_1 for the weak hypercharge $U(1)_Y$. The group generators are the weak isospin operators I_W^a (where $a = 1, 2, 3$) and the weak hypercharge operator Y_W . These abstract quantum numbers are related to the electric charge Q via the Gell-Mann-Nishijima formula:

$$Q = I_W^3 + \frac{Y_W}{2}. \quad (1.1)$$

Gauge invariance under G_{EW} requires the introduction of four vector fields transforming in the adjoint representation of the group. The $SU(2)_L$ interaction is mediated by the isotriplet $W_\mu^a(x)$, while the $U(1)_Y$ interaction is mediated by the singlet $B_\mu(x)$. The dynamics of these bosons are governed by their respective field strength tensors:

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu, \quad W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + g_2 \varepsilon^{abc} W_\mu^b W_\nu^c, \quad (1.2)$$

where ε^{abc} represents the totally antisymmetric structure constants of $SU(2)$. The coupling between the gauge bosons and matter fields is introduced via the covariant derivative³:

$$D_\mu = \partial_\mu - ig_2 \frac{\sigma^a}{2} W_\mu^a + ig_1 \frac{Y_W}{2} B_\mu, \quad (1.3)$$

where σ^a are the Pauli matrices. In the symmetric phase, the gauge bosons described above are massless. To generate mass without violating gauge invariance, the Higgs mechanism is employed to spontaneously break the symmetry $SU(2)_L \times U(1)_Y \rightarrow U(1)_{em}$. This is achieved by introducing a complex scalar doublet $\Phi(x)$ with weak hypercharge $Y_{W,\Phi} = \pm 1$. This hypercharge assignment ensures the doublet contains a neutral component capable of acquiring a non-vanishing vacuum expectation value (vev). The doublet is parameterized as:

$$\Phi(x) = \begin{pmatrix} \varphi^+(x) \\ \varphi^0(x) \end{pmatrix}. \quad (1.4)$$

Crucially, the assignment $Y_{W,\Phi} = +1$ permits gauge-invariant Yukawa couplings between the scalar field and fermions, providing the mechanism for fermion mass generation upon symmetry breaking.

The theory is chiral, treating left- and right-handed fermion fields differently. Left-handed fields transform as doublets under $SU(2)_L$. Denoting the lepton and quark doublets as L and Q respectively:

$$L_j'^L = \omega_- L_j' = \begin{pmatrix} \nu_j'^L \\ l_j'^L \end{pmatrix}, \quad Q_j'^L = \omega_- Q_j' = \begin{pmatrix} u_j'^L \\ d_j'^L \end{pmatrix}. \quad (1.5)$$

Conversely, right-handed fields transform as $SU(2)_L$ singlets:

$$l_j'^R = \omega_+ l_j', \quad u_j'^R = \omega_+ u_j', \quad d_j'^R = \omega_+ d_j'. \quad (1.6)$$

Here, the chirality projectors are defined as $\omega_\pm = \frac{1}{2}(1 \pm \gamma^5)$, and the index $j = 1, 2, 3$ denotes the generation. The labels ν , l , u , and d correspond to neutrinos, charged leptons, up- and down-type quarks,

³We adopt the conventions of [35–38]. Different sign conventions are used in the literature. Often the form $D_\mu = \partial_\mu + ig_2 \sigma^a / 2 W_\mu^a + ig_1 Y_W B_\mu$ is used [39, 40], which differs in the sign of the $SU(2)$ gauge fields W^a . In the field basis corresponding to charge and mass eigenstates this corresponds to a sign change in the $W^\pm$ and Z fields. Peskin and Schroeder [24] and Schwartz [26], on the other hand, use the convention $D_\mu = \partial_\mu - ig_2 \sigma^a / 2 W_\mu^a - ig_1 Y_W B_\mu$, which is related to our convention by a sign change in the $U(1)_Y$ gauge field B_μ and the photon field A_μ . Parameter relations are not affected by these transformations. However, the signs of Feynman rules and Green functions change.

respectively. Right-handed neutrinos are omitted from this discussion as they are not relevant for the collider phenomenology considered here. The prime notation indicates the interaction basis, where the covariant derivative is diagonal in flavor space; these states are not necessarily mass eigenstates. Due to the chiral nature of the gauge interactions, explicit fermion mass terms are forbidden in the Lagrangian. The complete classical Lagrangian $\mathcal{L}_{\text{class}}$ for the electroweak sector is given by:

$$\begin{aligned} \mathcal{L}_{\text{class}} = & -\frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W_{\mu\nu}^a W^{a,\mu\nu} \\ & + (D_\mu\Phi)^\dagger (D^\mu\Phi) + \mu^2 (\Phi^\dagger\Phi) - \frac{\lambda}{4} (\Phi^\dagger\Phi)^2 \\ & + \sum_i \left(\overline{L_i'^L} i \not{D} L_i'^L + \overline{Q_i'^L} i \not{D} Q_i'^L \right) \\ & + \sum_i \left(\overline{l_i'^R} i \not{D} l_i'^R + \overline{u_i'^R} i \not{D} u_i'^R + \overline{d_i'^R} i \not{D} d_i'^R \right) \\ & - \sum_{i,j} \left(\overline{L_i'^L} G_{ij}^l l_j'^R \Phi + \overline{Q_i'^L} G_{ij}^u u_j'^R \Phi^c + \overline{Q_i'^L} G_{ij}^d d_j'^R \Phi + \text{h.c.} \right), \end{aligned} \quad (1.7)$$

The first line describes the gauge boson kinetic terms. The second line details the Higgs sector, including its kinetic term, gauge interactions, and the scalar potential characterized by the mass parameter μ^2 and quartic coupling λ . The third and fourth lines represent the fermion kinetic terms and their gauge interactions. The final line contains the Yukawa interactions, parameterized by the coupling matrices G_{ij}^f . To couple to the up-type quarks, the charge-conjugate Higgs doublet is required, defined as $\Phi^c(x) = i\tau_2 \Phi^*(x) = (\varphi^{0*}(x), -\varphi^-(x))^T$ with τ_2 the second Pauli matrix.

In summary, the classical Lagrangian is fully determined by the gauge couplings g_1 and g_2 , the Higgs potential parameters μ^2 and λ , and the Yukawa matrices G_{ij}^f .

1.2.2 Lagrangian of the electroweak SM in the physical basis

Our analysis begins with the scalar potential of the Higgs sector:

$$V(\Phi) = -\mu^2 (\Phi^\dagger\Phi) + \frac{\lambda}{4} (\Phi^\dagger\Phi)^2. \quad (1.8)$$

The vacuum structure depends critically on the sign of the parameter μ^2 . If $\mu^2 < 0$, the minimum is located at $\Phi_0 = (0, 0)^T$, leaving the gauge symmetry unbroken and the gauge bosons massless. However, in the Standard Model, we assume $\mu^2 > 0$, which destabilizes the origin and shifts the minimum to a non-zero value satisfying:

$$(\Phi_0^\dagger\Phi_0) = \frac{2\mu^2}{\lambda} \neq 0. \quad (1.9)$$

Due to the $SU(2)$ invariance of the potential, there is an infinite set of degenerate vacua. Spontaneous symmetry breaking occurs when the physical system settles into one of these ground states. By exploiting the EW gauge freedom, we can, without loss of generality, align this vacuum expectation value (vev) along a specific direction:

$$\Phi_0 = \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix}. \quad (1.10)$$

where v is real and positive. This solution is unique up to a phase. We then expand the scalar field $\Phi(x)$ around this vev v :

$$\Phi(x) = \begin{pmatrix} \varphi^+(x) \\ \frac{1}{\sqrt{2}} [v + H(x) + i\chi(x)] \end{pmatrix}, \quad \varphi^-(x) = [\varphi^+(x)]^\dagger, \quad (1.11)$$

Here, $H(x)$ represents the physical, massive Higgs boson. The fields $\chi(x)$ and $\varphi^\pm(x)$ represent the massless would-be Goldstone bosons. In the unitary gauge, these unphysical degrees of freedom are reabsorbed by the gauge bosons, providing the longitudinal polarization states necessary for them to acquire mass. In a general gauge the longitudinal degrees of freedom of the gauge boson combine with the would-be Goldstone in a specific combination which is controlled by a Ward identity.

To derive the particle spectrum in the physical basis, we substitute the expansion Eq. 1.11 into the classical Lagrangian (Eq. 1.7) and diagonalize the resulting mass matrices. The physical masses are found to be:

$$\begin{aligned} m_W &= \frac{1}{2}vg_2, & m_Z &= \frac{1}{2}v\sqrt{g_1^2 + g_2^2}, & m_\gamma &= 0, \\ m_H &= \sqrt{2\mu^2} = v\sqrt{\frac{\lambda}{2}}, & m_{f,i} &= \frac{v}{\sqrt{2}} \sum_{k,m} U_{ik}^{f,L} G_{km}^f U_{mi}^{f,R\dagger}. \end{aligned} \quad (1.12)$$

The corresponding physical eigenstates relate to the gauge eigenstates as follows:

$$\begin{aligned} W_\mu^\pm &= \frac{1}{\sqrt{2}} (W_\mu^1 \mp iW_\mu^2), & \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix} &= \begin{pmatrix} c_W & s_W \\ -s_W & c_W \end{pmatrix} \begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix}, \\ f_i^L &= \sum_k U_{ik}^{f,L} f_k'^L, & f_i^R &= \sum_k U_{ik}^{f,R} f_k'^R, \end{aligned} \quad (1.13)$$

Here, we have introduced the weak mixing angle θ_W , where $c_W \equiv \cos \theta_W = g_2/\sqrt{g_1^2 + g_2^2} = m_W/m_Z$ and $s_W \equiv \sin \theta_W$. The electromagnetic coupling constant e is related to the original gauge couplings by:

$$e = \frac{g_1 g_2}{\sqrt{g_1^2 + g_2^2}} = g_2 s_W = g_1 c_W = \sqrt{4\pi\alpha}. \quad (1.14)$$

After electroweak symmetry breaking, the Yukawa interactions generate explicit mass matrices for the fermions. To transition to the physical mass basis, these matrices are diagonalized via bi-unitary transformations involving the matrices $U^{f,L}$ and $U^{f,R}$. Crucially, while $SU(2)_L$ gauge invariance requires the upper and lower components of left-handed doublets to transform identically before electroweak symmetry breaking, they can be rotated independently once the symmetry is broken. In the neutral current interactions, the unitarity of these rotation matrices ensures that flavor-changing neutral currents are absent at the tree level. However, in the charged current interactions, the independent rotations of the left-handed up-type and down-type quarks do not align, resulting in the Cabibbo-Kobayashi-Maskawa (CKM) mixing matrix [41, 42]:

$$V = U^{u,L} U^{d,L\dagger}. \quad (1.15)$$

For the lepton sector, because neutrinos remain massless in this minimal framework, their basis is degenerate. We can therefore utilize this freedom to choose $U^{\nu,L} = U^{l,L}$, perfectly aligning the flavor and mass bases and rendering the neutrino interactions flavor-diagonal.

Expressing the classical Lagrangian in terms of these physical fields yields the complete electroweak Lagrangian:

$$\begin{aligned}
\mathcal{L}_{\text{class}} = & \sum_{f=l,\nu,u,d} \sum_i [\bar{f}_i(i\not{\partial} - m_{f,i})f_i - eQ_f \bar{f}_i \gamma^\mu f_i A_\mu] \\
& + \sum_{f=l,\nu,u,d} \sum_i \frac{e}{s_W c_W} [I_{W,f}^3 \bar{f}_i^L \gamma^\mu f_i^L - s_W^2 Q_f \bar{f}_i \gamma^\mu f_i] Z_\mu \\
& + \sum_{i,j} \frac{e}{\sqrt{2}s_W} [\bar{u}_i^L \gamma^\mu V_{ij} d_j^L W_\mu^+ + \bar{d}_i^L \gamma^\mu V_{ij}^\dagger u_j^L W_\mu^-] \\
& + \sum_i \frac{e}{\sqrt{2}s_W} [\bar{\nu}_i^L \gamma^\mu l_i^L W_\mu^+ + \bar{l}_i^L \gamma^\mu \nu_i^L W_\mu^-] \\
& - \frac{1}{4} |\partial_\mu A_\nu - \partial_\nu A_\mu - ie(W_\mu^- W_\nu^+ - W_\nu^- W_\mu^+)|^2 \\
& - \frac{1}{4} \left| \partial_\mu Z_\nu - \partial_\nu Z_\mu + ie \frac{c_W}{s_W} (W_\mu^- W_\nu^+ - W_\nu^- W_\mu^+) \right|^2 \\
& - \frac{1}{2} \left| \partial_\mu W_\nu^+ - \partial_\nu W_\mu^+ - ie(W_\mu^+ A_\nu - W_\nu^+ A_\mu) \right. \\
& \quad \left. + ie \frac{c_W}{s_W} (W_\mu^+ Z_\nu - W_\nu^+ Z_\mu) \right|^2 \\
& + \frac{1}{2} \left| \partial_\mu (H + i\chi) - i \frac{e}{s_W} W_\mu^- \varphi^+ - im_Z Z_\mu + i \frac{e}{2c_W s_W} Z_\mu (H + i\chi) \right|^2 \\
& + \left| \partial_\mu \varphi^+ + ie A_\mu \varphi^+ - ie \frac{c_W^2 - s_W^2}{2c_W s_W} Z_\mu \varphi^+ - im_W W_\mu^+ \right. \\
& \quad \left. - i \frac{e}{2s_W} W_\mu^+ (H + i\chi) \right|^2 \\
& - \frac{1}{2} m_H^2 H^2 - \frac{m_H^2}{v} H \left(\varphi^- \varphi^+ + \frac{1}{2} |H + i\chi|^2 \right) \\
& - \frac{m_H^2}{2v^2} \left(\varphi^- \varphi^+ + \frac{1}{2} |H + i\chi|^2 \right)^2 \\
& - \sum_{f=l,u,d} \sum_i \frac{m_{f,i}}{v} (\bar{f}_i f_i H - 2I_{W,f}^3 \bar{f}_i \gamma^5 f_i \chi) \\
& + \sum_{i,j} \frac{\sqrt{2}}{v} \left[m_{u,i} \bar{u}_i^R V_{ij} d_j^L \varphi^+ + m_{u,j} \bar{d}_i^L V_{ij}^\dagger u_j^R \varphi^- \right. \\
& \quad \left. - m_{d,j} \bar{u}_i^L V_{ij} d_j^R \varphi^+ - m_{d,i} \bar{d}_i^R V_{ij}^\dagger u_j^L \varphi^- \right] \\
& - \sum_i \frac{\sqrt{2}}{v} m_{l,i} (\bar{\nu}_i^L l_i^R \varphi^+ + \bar{l}_i^R \nu_i^L \varphi^-). \tag{1.16}
\end{aligned}$$

1.2.3 Quantization

To perform the quantization of the theory, a specific gauge must be selected to remove the redundancy of the gauge freedom. This is achieved by supplementing the Lagrangian with a gauge-fixing term, defined

as:

$$\mathcal{L}_{\text{fix}} = -\frac{1}{2\xi_A}(C^A)^2 - \frac{1}{2\xi_Z}(C^Z)^2 - \frac{1}{2\xi_W}C^+C^- \quad (1.17)$$

where ξ_a represents the arbitrary gauge parameters and C^a denotes the linear gauge-fixing operators. In the specific R_ξ gauges (or 't Hooft gauges) adopted here, these operators take the form:

$$C^A = \partial^\mu A_\mu, \quad C^Z = \partial^\mu Z_\mu - m_Z \xi_Z \chi, \quad C^\pm = \partial^\mu W_\mu^\pm \mp i m_W \xi_W \varphi^\pm. \quad (1.18)$$

The inclusion of $\mathcal{L}_{\text{fix}}$ allows for the inversion of the quadratic kinetic terms for the gauge fields, yielding the general vector boson propagator:

$$i\Delta_{\mu\nu}^{V^a}(q) = \frac{i}{q^2 - m_a^2} \left(-g_{\mu\nu} + \frac{q_\mu q_\nu}{q^2} \right) - \frac{i\xi_a}{q^2 - \xi_a m_a^2} \frac{q_\mu q_\nu}{q^2}, \quad (1.19)$$

where a corresponds to A, Z , or $W^\pm$. It is important to note that the transverse component of the propagator, which describes the physical degrees of freedom, is independent of ξ_a and thus gauge-invariant. Conversely, the longitudinal component is gauge-dependent. Furthermore, the chosen gauge-fixing operators induce quadratic terms for the Goldstone bosons, effectively generating gauge-dependent masses:

$$-\xi_W m_W^2 \varphi^+ \varphi^- - \frac{1}{2} \xi_Z m_Z^2 \chi^2. \quad (1.20)$$

The gauge parameters ξ_a may be chosen freely to simplify specific calculations. A convenient choice is the Feynman-'t Hooft gauge, defined by $\xi_a = 1$. In this gauge, the vector propagator contains only the metric term $-g_{\mu\nu}$, and the unphysical Goldstone bosons acquire masses identical to their corresponding gauge bosons, i.e. $m_{\text{Goldstone}} = m_{\text{gauge}}$. This cancellation of poles makes the 't Hooft-Feynman gauge particularly advantageous for higher-order perturbative calculations. Alternatively, one may consider the unitary gauge, defined by the limit $\xi_a \rightarrow \infty$. In this regime, the Goldstone masses diverge, causing these unphysical degrees of freedom to decouple entirely from the spectrum, thereby making the physical content of the theory more transparent at the expense of more complex propagators.

The final requirement for a consistent quantization is the restoration of unitarity via the Faddeev-Popov procedure. This entails the introduction of ghost fields, described by the Lagrangian:

$$\mathcal{L}_{\text{FP}} = - \int d^4y \bar{u}^a(x) \frac{\delta C^a(x)}{\delta \theta^b(y)} u^b(y), \quad (1.21)$$

where $\delta C^a(x)/\delta \theta^b(y)$ represents the variation of the gauge-fixing condition under an infinitesimal gauge transformation, and u^a (with $a = A, Z, \pm$) are the anticommuting scalar ghost fields.

The complete Lagrangian for the electroweak sector of the Standard Model is thus the sum of the classical, gauge-fixing, and ghost terms:

$$\mathcal{L}_{\text{EW}} = \mathcal{L}_{\text{class}} + \mathcal{L}_{\text{fix}} + \mathcal{L}_{\text{FP}}. \quad (1.22)$$

1.2.4 Inclusion of strong interaction

The strong interaction is described by QCD, a non-Abelian gauge theory formulated in a manner analogous to the electroweak sector. The underlying gauge group is $SU(3)_c$, which possesses eight generators corresponding to eight massless gluon fields, denoted by G_μ^A (with $A = 1, \dots, 8$). The dynamics of

these fields is governed by the gluon field strength tensor:

$$G_{\mu\nu}^A = \partial_\mu G_\nu^A - \partial_\nu G_\mu^A + g_s f^{ABC} G_\mu^B G_\nu^C, \quad (1.23)$$

where g_s represents the strong coupling constant and f^{ABC} are the structure constants of the $SU(3)$ group. To incorporate QCD into the Standard Model, the classical Lagrangian from Eq. 1.7 is augmented by the kinetic term for the gluon fields:

$$\mathcal{L}_{\text{class}} \rightarrow \mathcal{L}_{\text{class}} - \frac{1}{4} G_{\mu\nu}^A G^{A,\mu\nu}. \quad (1.24)$$

The interactions between the gauge bosons and matter fields are established by generalizing the covariant derivative in Eq. 1.3 to include the color sector:

$$D_\mu = \partial_\mu - ig_s T^A G_\mu^A - ig_2 I_W^a W_\mu^a + ig_1 \frac{Y_W}{2} B_\mu. \quad (1.25)$$

Here, T^A denotes the generators of $SU(3)$. The quarks transform under the fundamental representation of the color group, for which the generators are given by $T^A = \lambda^A/2$, with λ^A being the Gell-Mann matrices. In contrast, the leptons, the Higgs doublet, and the electroweak gauge bosons are singlets under $SU(3)_c$ (i.e., $T^A = 0$) and do not participate in strong interactions. Conversely, the gluon fields themselves carry no weak isospin or hypercharge quantum numbers.

1.3 Background-field method

As discussed in Sec. 1.2.3, the quantization of the Standard Model within the conventional formalism necessitates a gauge-fixing term $\mathcal{L}_{\text{fix}}$. This term explicitly breaks gauge invariance at intermediate stages of the calculation. The symmetry is only restored in the final results upon projection onto the physical degrees of freedom. An alternative approach that circumvents this issue is the Background-Field Method (BFM). Originally introduced in the context of quantum gravity to handle the complexities of non-Abelian gauge symmetries [43], the method was subsequently generalized to the QCD and EW sectors of the SM [44–48].

The central premise of the BFM is the decomposition of the fields in the classical Lagrangian into a classical background configuration, $\hat{\Psi}$, and quantum fluctuations, Ψ :

$$\mathcal{L}_{\text{class}}(\Psi) \rightarrow \mathcal{L}_{\text{class}}(\Psi + \hat{\Psi}). \quad (1.26)$$

Quantization proceeds by integrating over the quantum fields Ψ while treating $\hat{\Psi}$ as an external source. Crucially, one adopts a specific gauge-fixing term that breaks the gauge invariance of the Lagrangian under infinitesimal gauge transformation of the quantum fields Ψ , while preserving the invariance under transformations of the background fields $\hat{\Psi}$. This gauge-fixing term is given by:

$$\mathcal{L}_{\text{fix,BFM}} = -\frac{1}{2\xi_Q} [(G^A)^2 + (G^Z)^2 + 2G^+ G^-], \quad (1.27)$$

to work with scalar ones. In order to do so we decompose them accordingly

$$\begin{aligned}
\Sigma_{\mu\nu}^{VV}(p^2) &= \left(-g_{\mu\nu} + \frac{p_\mu p_\nu}{p^2} \right) \Sigma_T^{VV}(p^2) - \frac{p_\mu p_\nu}{p^2} \Sigma_L^{VV}(p^2), \\
\Sigma_\mu^{VS}(p^2) &= p_\mu \Sigma^{VS}(p^2), \\
\Sigma^{SS}(p^2) &= \Sigma^{SS}(p^2), \\
\Sigma^{ff}(p^2) &= \not{p} \omega_- \Sigma_L^{ff}(p^2) + \not{p} \omega_+ \Sigma_R^{ff}(p^2) + m_f \Sigma_S^{ff}(p^2).
\end{aligned} \tag{1.31}$$

In terms of the scalar self-energies defined in Eq. 1.31, the Ward identities take the form:

$$\begin{aligned}
\Sigma_L^{\hat{A}\hat{A}}(p^2) &= 0, & \Sigma_T^{\hat{A}\hat{A}}(0) &= 0, \\
\Sigma_L^{\hat{A}\hat{Z}}(p^2) &= 0, & \Sigma_T^{\hat{A}\hat{Z}}(0) &= 0, \\
\Sigma^{\hat{A}\hat{\chi}}(p^2) &= 0, & \Sigma^{\hat{A}\hat{H}}(p^2) &= 0, \\
\Sigma_L^{\hat{Z}\hat{Z}}(p^2) - im_Z \Sigma^{\hat{Z}\hat{\chi}}(p^2) &= 0, \\
p^2 \Sigma^{\hat{Z}\hat{\chi}}(p^2) - im_Z \Sigma^{\hat{\chi}\hat{\chi}}(p^2) + i \frac{e}{2c_W s_W} T^{\hat{H}} &= 0, \\
\Sigma_L^{\hat{W}^\pm \hat{W}^\mp}(p^2) \mp m_W \Sigma^{\hat{W}^\mp \hat{\phi}^\pm}(p^2) &= 0, \\
p^2 \Sigma^{\hat{W}^\pm \hat{\phi}^\mp}(p^2) \mp m_W \Sigma^{\hat{\phi}^\pm \hat{\phi}^\mp}(p^2) \pm \frac{e}{2s_W} T^{\hat{H}} &= 0.
\end{aligned} \tag{1.32}$$

These relations, in the BFM formalism, hold at any order in perturbation theory. Note that the vanishing of the photon-Z-boson mixing at zero momentum is explicitly enforced through a renormalization condition in the usual on-shell scheme, while in the BFM it is automatically fulfilled as a consequence of gauge invariance. We can write down Ward identities also for three-point functions. For example, the vertex function $\Gamma_\mu^{\hat{A}\bar{f}f}$ obeys:

$$k^\mu \Gamma_\mu^{\hat{A}\bar{f}f}(k, \bar{p}, p) = -eQ_f \left[\Gamma^{\bar{f}f}(\bar{p}) - \Gamma^{\bar{f}f}(-p) \right]. \tag{1.33}$$

This QED-like Ward identity implies that the renormalization of the electric charge depends only on the photon self-energy, leading to the universality of the electric charge.

1.5 Problems of the Standard Model

The SM is one of the most rigorously tested and experimentally validated theoretical frameworks in modern physics. During the years, it has been extensively explored at collider such as the Large Hadron Collider (LHC) in Geneva, Switzerland, and no significant evidence of BSM physics has been found. Among its greater successes there is the prediction of the Higgs boson, which has been produced and observed in 2012, and the measurement of anomalous magnetic dipole moment of the electron, which has been tested against theoretical predictions up to the fifth perturbative order! Despite its incredible achievements, however, there are some problems which cannot be explained within the SM, and, hence, require a more general theory, yet to be found. Here is a non-comprehensive list of unsolved problems.

- **Neutrino masses:** in the original, minimal formulation of the SM, fermions gain mass by interacting with the Higgs field, a mechanism that requires both left-handed and right-handed chiral states. Because this minimal model includes only left-handed neutrinos, they were initially assumed to be massless. However, the definitive experimental discovery of neutrino flavor oscillations proved

that neutrinos do possess a non-zero mass [50–54]. While it is theoretically possible to extend the SM by adding right-handed neutrinos, which would be sterile under the SM gauge group, to generate standard Dirac masses, this naive inclusion introduces a naturalness problem. The observed neutrino masses are extraordinarily small compared to all other fermions. This extreme hierarchy suggests that neutrino masses are not generated by the standard Higgs mechanism alone, but rather through a BSM dynamics operating at energy scales far above the EW scale, such as the see-saw mechanism. For a review on the problem see [55, 56].

- **Strong CP problem:** in the QCD part of the Lagrangian, the $SU(3)_c$ symmetry allows the presence of a term

$$\mathcal{L}_\theta = -\theta_{\text{QCD}} \frac{g_s^2}{16\pi^2} \tilde{G}_{\mu\nu}^a G^{a,\mu\nu} \quad (1.34)$$

where the dual-field strength is defined through the Levi-Civita tensor as

$$\tilde{G}_{\mu\nu}^a = -\frac{1}{2} \epsilon_{\mu\nu\rho\sigma} G^{a,\rho\sigma}, \quad (1.35)$$

with $\epsilon^{0123} = +1$. This Lagrangian violates CP symmetry and would give rise to experimentally measurable quantities such as a neutron electric dipole moment. Current experimental limits on this quantity, however, translate to an upper limit on $|\theta_{\text{QCD}}| < 10^{-10}$ [57–59]. This is interesting since θ_{QCD} , in principle, could take any value from 0 to 2π . The fine-tuning of θ_{QCD} cannot be explained within the SM and requires a solution beyond it, for example, the existence of axions [26, 60].

- **Dark matter:** observations of galactic rotation curves suggest that only 20% of the matter content of the universe can be described by the SM. In particular, stars in the outer galactic regions exhibit constant rotational velocities, rather than the expected Keplerian decline $v \propto r^{1/2}$ [61, 62]. This phenomenon violates Newtonian dynamics unless galaxies are embedded within a massive halo of invisible matter. This missing component, termed Dark Matter, must be stable, electrically neutral, and non-baryonic. While SM neutrinos possess these properties, their minuscule masses dictate that they would constitute highly relativistic dark matter. This is incompatible with observations of cosmological large-scale structure formation and insufficient to account for the required relic density, leaving the SM without a viable candidate.
- **Matter-Antimatter asymmetry:** The observable universe exhibits a profound baryon asymmetry, that is a near-total dominance of matter over antimatter. Assuming a symmetric initial state following cosmic inflation, a universe governed solely by the SM should have produced and maintained equal amounts of matter and antimatter, ultimately leading to total annihilation. To generate this asymmetry dynamically, the universe must satisfy the three Sakharov conditions: baryon number violation, C and CP violation, and a departure from thermal equilibrium. While the SM technically contains all three ingredients, they are quantitatively insufficient by roughly 10 orders of magnitude [63, 64].
- **Gravity:** As already mentioned, the SM does not include gravity. While General Relativity can be successfully quantized as an EFT to describe graviton-mediated interactions at currently accessible energies, this approach is perturbatively non-renormalizable. Because the gravitational coupling has a negative mass dimension, the EFT framework loses predictive power and produces uncontrollable infinities as interaction energies approach the Planck scale, $M_{\text{Pl}} \sim 10^{19}$ GeV. Consequently,

the SM and this gravitational EFT fail completely in environments where extreme macroscopic masses and microscopic quantum scales exist simultaneously, such as the singularity of a black hole or the earliest moments of the Big Bang.

In conclusion, we have both theoretically and experimental reasons to believe that the SM is just a low-energy approximation of a more fundamental theory. During the years many BSM theories have been proposed, however none of them has yet found an experimental confirmation.

1.6 Standard Model Effective Field Theory

Despite the comprehensive experimental program at the LHC, direct evidence for physics Beyond the Standard Model (BSM) remains elusive in the TeV energy regime. This suggests that if new physics exists, it likely resides at a mass scale Λ significantly higher than the electroweak scale ($v \approx 246$ GeV). In this "mass gap" scenario, heavy BSM particles cannot be produced on-shell, but their virtual effects may manifest as small deviations in the high-energy tails of collider observables.

The Standard Model Effective Field Theory (SMEFT) provides a robust, model-independent framework to parameterize these potential deviations. The underlying premise is based on the decoupling theorem: one assumes a UV-complete theory exists at the scale Λ , and by integrating out the heavy degrees of freedom, one derives an effective Lagrangian valid at energies $E \ll \Lambda$. This effective Lagrangian is organized as a systematic expansion in inverse powers of the new physics scale:

$$\mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}} + \frac{1}{\Lambda} \sum_k c_k^{(5)} \mathcal{O}_k^{(5)} + \frac{1}{\Lambda^2} \sum_k c_k^{(6)} \mathcal{O}_k^{(6)} + \mathcal{O}\left(\frac{1}{\Lambda^3}\right), \quad (1.36)$$

where $\mathcal{L}_{\text{SM}}$ represents the renormalizable Standard Model Lagrangian in Eq. 1.22. The term $\mathcal{O}_k^{(D)}$ denotes a complete set of gauge-invariant operators of mass dimension D , constructed solely from SM fields. The dimensionless Wilson coefficients $c_k^{(D)}$ encode the effects of the heavy physics. In particular, they depend on the couplings and masses of the integrated-out heavy states.

The leading correction to the SM arises at dimension-5. There is only one class of such operators (up to Hermitian conjugation and flavor structure), known as the Weinberg operator:

$$\mathcal{O}^{(5)} = (\bar{L}^c \tilde{\Phi})^* (\tilde{\Phi}^\dagger L). \quad (1.37)$$

This operator violates lepton number by two units ($\Delta L = 2$) and, upon symmetry breaking, generates Majorana masses for the neutrinos. The extreme smallness of neutrino masses suggests that the scale Λ associated with this operator is very high (near the GUT scale), justifying the neglect of $D = 5$ effects in most collider physics phenomenology.

For this reason, the primary focus of LHC phenomenology lies at dimension-6. While the dimension-6 basis contains four operator structures that violate baryon and lepton number (which expand into hundreds of independent parameters when flavor indices are included), the vast majority of phenomenological studies focus on the baryon and lepton conserving sector. These operators contribute to Higgs physics, top quark properties, and electroweak precision observables. The choice of operators $\mathcal{O}_k^{(6)}$ is not unique. Operators can be related via field redefinitions and Fierz identities. These redundancies essentially allow one to trade one operator for a linear combination of others without affecting on-shell S -matrix elements.

The description of BSM physics based on an effective Lagrangian in terms of the SM fields was pioneered by Buchmuller and Wyler [65], who provided a list of operators of dimensions 5 and 6 in

the linear parametrization of the Higgs sector with a Higgs doublet. While a complete minimal basis which also eliminates these redundancies is the so-called *Warsaw Basis* [40]. This basis consists of 59 distinct operator classes which conserves baryon and lepton number. When including the flavor structure of the three fermion generations, this expands to 2499 independent Wilson coefficients. Other bases were proposed, and different authors prefer to use different sets of operator bases, motivated by different expectations for the form of new physics [66–72].

A crucial theoretical feature of the SMEFT is that it is renormalizable order-by-order in the E/Λ expansion. Although the inclusion of higher-dimensional operators with negative mass-dimension couplings renders the theory non-renormalizable in the traditional sense, SMEFT is mathematically consistent as a quantum field theory at any truncated order. For instance, ultraviolet divergences arising from loop corrections that involve exactly one insertion of a dimension-6 operator can be entirely absorbed by the counterterms of the dimension-6 operators themselves. This procedure naturally induces operator mixing, meaning that a Wilson coefficient evaluated at a high scale Λ will "run" and mix with other operators as it is evolved down to the data scale. This scale dependence is rigorously governed by Renormalization Group Equations (RGEs) driven by the anomalous dimension matrix of the SMEFT. Ultimately, order-by-order renormalizability is what permits consistent next-to-leading order calculations within the SMEFT, ensuring that experimental measurements can be systematically compared against high-precision theoretical predictions.

Handling 2499 free parameters is experimentally intractable. To make the global analysis of data feasible, one must impose specific flavor symmetries. A common assumption is Minimal Flavor Violation (MFV) or a $U(3)^5$ flavor symmetry, where the Yukawa matrices remain the only source of flavor breaking. Under the assumption that the BSM sector is flavor-blind (or only breaks flavor universality in specific sectors like the third generation), the number of relevant parameters reduces drastically to $\mathcal{O}(10 - 20)$ for specific processes, such as Higgs or top production.

The validity of the SMEFT description relies on the condition that the typical energy transfer of the process E satisfies $E < \Lambda$. If the interaction energy approaches the cut-off scale ($E \approx \Lambda$), provided that Λ is a physical quantity, e.g. the mass of a particle, the expansion in powers of E/Λ breaks down, and the full UV-complete theory is required. However, ensuring kinematic validity is not sufficient. One must also impose strict theoretical constraints, particularly perturbative unitarity. Because dimension-6 operators induce scattering amplitudes that grow with energy ($\mathcal{M} \propto c_k E^2/\Lambda^2$), this growth will eventually violate unitarity at some critical scale. Theoretical consistency therefore demands that the Wilson coefficients remain bounded, typically $c_k < 4\pi$, to ensure the EFT remains unitary and perturbative at the energies probed. In experimental analyses, it is thus crucial to verify two things: first, that the extracted limits on c_k/Λ^2 imply a Λ safely above the kinematic reach of the experiment and, second, that the resulting coefficients respect these theoretical unitarity bounds.

In summary, the SMEFT serves as a bridge between data and theory: experimentalists constrain the ratio c_k/Λ^2 , and theorists compute these coefficients by matching specific UV-complete models to the EFT at the scale Λ . This allows for a global stress-test of the Standard Model without committing to a specific BSM hypothesis.

Amplitudes

“The word adventure has gotten overused. For me, when everything goes wrong, that’s when adventure starts.”

Yvon Chouinard, *rock climber, alpinist, founder of “Patagonia”.*

In this Chapter we will introduce scattering amplitudes, which will constitute the main focus of this thesis, and we will discuss general problems related to their calculation, namely UV renormalization, the complex mass scheme and the treatment of γ_5 . After a review on general aspects of UV renormalization, we expand pioneering work at two-loop mixed QCD-EW [73] and two-loop EW [74–79] to systematically cover all required counterterms necessary for any process at two-loop mixed and two-loop EW level. These findings have been detailed in [8]. Subsequently, we review general aspects of the Complex Mass Scheme and we present two standard approaches for treating γ_5 in d dimensions: the ’t Hooft-Veltman-Breitenlohner-Maison (HVBM) [80–83] prescription and Kreimer’s [84] prescription. Following this, we detail the established methodology for computing d -dimensional Feynman integrals containing the four-dimensional scalar numerators that naturally emerge when applying Kreimer’s method. The theoretical aspects discussed in this subsection are drawn entirely from the literature. In the computation of the two-loop amplitudes that we will be presented in Chapter 5, we have consistently used Kreimer’s approach, which we implemented into our package ABISS (see Section 4.4.1).

The central objects of study in high-energy particle physics are the scattering amplitudes. They serve as the fundamental link between the abstract formalism of QFT, defined by a Lagrangian $\mathcal{L}$, and the measurable quantities observed in experiments, such as scattering cross-sections and decay rates.

Physically, the squared amplitudes $|\mathcal{M}|^2$ encode the probability for a specific quantum mechanical transition, i.e. a set of initial-state particles $|i\rangle$ interacting to produce a set of final-state particles $|f\rangle$. Mathematically, the amplitude is a matrix element of the scattering operator which, for a fixed set of states, it evaluates to a complex-valued function of the external kinematics (momenta, spins, and helicities) and the internal parameters of the theory (couplings, masses).

While the underlying quantum fields fluctuate continuously, scattering experiments typically observe particles at asymptotic times ($t \rightarrow \pm\infty$), where they are sufficiently separated to be treated as non-interacting. The amplitude bridges these asymptotic states through the interaction region.

The transition from an initial state $|i\rangle$ at $t = -\infty$ to a final state $|f\rangle$ at $t = +\infty$ is governed by the S-matrix operator $\hat{S}$. The probability amplitude for this process is the matrix element:

$$\langle f | \hat{S} | i \rangle. \quad (2.1)$$

Since particles often pass through each other without interacting, it is convenient to separate the S-matrix into the identity operator, meaning that there is no interaction, and the transition matrix $\hat{T}$:

$$\hat{S} = 1 + i\hat{T} \quad (2.2)$$

Momentum conservation is a universal symmetry in relativistic QFT. We therefore extract the momentum-conserving Dirac delta function to define the invariant scattering amplitude $\mathcal{M}$:

$$\langle f | i\hat{T} | i \rangle = (2\pi)^4 \delta^{(4)} \left(\sum_f p_f - \sum_i p_i \right) i\mathcal{M}(i \rightarrow f). \quad (2.3)$$

How do we compute $\mathcal{M}$ from the field operators in the Lagrangian? The connection is provided by the Lehmann-Symanzik-Zimmermann (LSZ) reduction formula.

The LSZ formula relates the S-matrix elements to the n -point Green's functions (time-ordered correlation functions) of the theory. It states that the scattering amplitude is obtained by taking the Green's function, amputating the external propagators, that is multiplying by the inverse propagator, and forcing the external momenta to be on-shell.

For a process involving n scalar particles, the LSZ formula reads:

$$\langle f | \hat{S} | i \rangle \sim \lim_{p_k^2 \rightarrow m_k^2} \left[\prod_{k=1}^n (p_k^2 - m_k^2) \right] G^{(n)}(p_1, \dots, p_n) \quad (2.4)$$

This procedure ensures that $\mathcal{M}$ contains only the physics of the interaction itself, stripping away the trivial propagation of the external particles before and after the collision.

Calculating the full Green's function non-perturbatively is generally impossible. Instead, we expand the interaction part of the Lagrangian, $\mathcal{L}_{\text{int}}$, in a power series of the coupling constants, α_s or α . This gives rise to the Feynman diagram expansion. A Feynman diagram is a graphical representation of a specific term in this perturbative series. It translates complex mathematical expressions into topological graphs built from a small set of Feynman Rules.

Feynman diagrams are topologically classified by the number of closed loops they contain, a distinction that dictates both their physical interpretation and mathematical complexity. At the lowest order, tree-level diagrams, i.e. connected graphs with no cycles, represent the leading contribution to the scattering process, where the flow of momentum is uniquely determined by the external kinematics. However, higher-order corrections introduce loop diagrams, where internal propagators form closed circuits. Physically, loops encode quantum fluctuations where virtual particles are created and annihilated within the interaction. Mathematically, they present a significantly steeper challenge: because the momentum circulating within the loop is not constrained by external momentum conservation, it must be integrated over all possible values in d -dimensional Minkowski space. These loop integrals are frequently divergent in the ultraviolet (UV) or infrared (IR) limits, necessitating regularization schemes, such as Dimensional Regularization, to isolate singularities and renormalization to absorb them into physical parameters. Furthermore, multi-loop amplitudes involve complex Feynman integrals, whose calculation will be the focus of Chapter 3.

The total amplitude $\mathcal{M}$ is constructed as the sum of these individual diagrammatic contributions. Consequently, the squared amplitude $|\mathcal{M}|^2 = \mathcal{M}^\dagger \mathcal{M}$, which is related to physical observables, as we will see in details in Chapter 4, involves not only the square of each diagram but also the interference terms between them, which must be summed over the color and spin states of the external particles to yield the final, fully contracted scalar quantity ready for phase-space integration.

2.1 UV renormalization

Renormalization in the Standard Model serves as the essential mathematical procedure that renders the theory predictive. It achieves this by systematically absorbing the UV singularities that arise in higher-order perturbative calculations into the bare parameters and fields, thereby allowing all theoretical predictions to be expressed in terms of finite, physically measurable quantities. Because the SM is a renormalizable non-abelian gauge theory with spontaneous symmetry breaking [85–89], these divergences can be systematically absorbed into renormalization constants that redefine the input parameters and fields. The general process requires selecting a specific renormalization scheme to fix these constants. While it is possible to renormalize the parameters in the symmetric basis of the original Lagrangian, practical calculations are most often performed using the On-Shell (OS) scheme [35–37, 90–93]. In this framework, the renormalization conditions are imposed directly on the physical basis, ensuring that the renormalized masses of gauge bosons, the Higgs, and fermions coincide with their physical masses as defined by the poles of their propagators, while the electric charge is fixed in the Thomson limit of low-energy scattering. This strategy not only assigns clear physical meaning to the parameters but also ensures that the calculation of S -matrix elements yields finite, testable predictions.

2.1.1 Renormalization of parameters and fields

We perform the renormalization for the parameters and fields of the physical basis introduced in Eq. 1.22. The bare parameters are split into renormalized parameters and counterterms as follows:

$$\begin{aligned} m_{W,0}^2 &= m_W^2 + \delta m_W^2, & m_{Z,0}^2 &= m_Z^2 + \delta m_Z^2, & m_{H,0}^2 &= m_H^2 + \delta m_H^2 \\ m_{f,i,0} &= m_{f,i} + \delta m_{f,i}, & e_0 &= e + \delta e, \end{aligned} \quad (2.5)$$

here m_W , m_Z , m_H and $m_{f,i}$ indicate the renormalized masses of the corresponding particles, while with the subscript 0 we denote the bare quantities. Similarly, we introduce renormalization constant, $Z = 1 + \delta Z$, for the fields via renormalization matrices

$$\begin{aligned} \hat{W}_0^\pm &= Z_W^{1/2} \hat{W}^\pm = (1 + \delta Z_W)^{1/2} \hat{W}^\pm \\ \begin{pmatrix} \hat{Z}_0 \\ \hat{A}_0 \end{pmatrix} &= \begin{pmatrix} Z_{ZZ}^{1/2} & (Z_{ZA}^{1/2} - 1) \\ (Z_{AZ}^{1/2} - 1) & Z_{AA}^{1/2} \end{pmatrix} \begin{pmatrix} \hat{Z} \\ \hat{A} \end{pmatrix} \\ \hat{H}_0 &= Z_H^{1/2} \hat{H} = (1 + \delta Z_H)^{1/2} \hat{H} \\ f_{0,i}^{L/R} &= (Z_f^{L/R})^{1/2} f_i^{L/R} = (1 + \delta Z_f^{L/R})^{1/2} f_{0,i}^{L/R}. \end{aligned} \quad (2.6)$$

Because the Z factor of every internal propagator algebraically cancels against the $Z^{-1/2}$ factors at its connecting vertices, the overall wave-function renormalization of internal ghosts and Goldstone fields vanishes entirely. As a result, only the uncanceled Z factors of the external physical states remain, rendering the renormalization constants of the ghosts and Goldstone sector irrelevant to the final amplitude calculation. The renormalization constants δX are intended as a double expansion in either the fine

structure constant, α , or the strong coupling, α_s :

$$\begin{aligned}
 \delta X &= \sum_{i,j} \alpha_s^i \alpha^j \delta X^{(i,j)} = \\
 &= \alpha \delta X^{(0,1)} + \alpha^2 \delta X^{(0,2)} + \dots \\
 &\quad + \alpha_s \delta X^{(1,0)} + \alpha_s \alpha \delta X^{(1,1)} + \dots \\
 &\quad + \alpha_s^2 \delta X^{(2,0)} + \dots
 \end{aligned} \tag{2.7}$$

2.1.2 Renormalization conditions

The renormalization constants are fixed by imposing renormalization conditions. These consist of two sets: the conditions that define the renormalized physical parameters and those that define the renormalized fields corresponding to physical particle states. While only the first set is relevant for the calculation of observables (S -matrix elements), a clever choice of the second set allows us not only to eliminate the explicit wave-function renormalization of the external particles, but also to simplify the explicit form of the renormalization conditions for the physical parameters. The renormalized mass parameters of the particles are fixed by the requirement that they are equal to the physical masses, which are defined via the locations of the poles of the corresponding propagators. These poles are equivalent to the zeros of the inverse connected 2-point functions projected onto physical states, i.e. on polarization vectors $\varepsilon^\mu(k)$ for gauge bosons and on spinors $u(p)$ and $v(p)$ for fermions. These can be considerably simplified by requiring the OS conditions for the field renormalization matrices in addition. The latter conditions require that close to its pole each propagator is given by its lowest-order expression with the bare mass replaced by the renormalized mass. As a consequence, an OS particle does not mix with others, and its propagator has residue one, i.e. the quanta of the renormalized fields are canonically normalized mass eigenstates. The renormalization conditions for the renormalized 2-point functions Γ_R for OS fields for physical external states, thus take the form [36]:

$$\begin{aligned}
 \Gamma_{R,\mu\nu}^{\hat{V}\hat{V}}(k, m_V) \varepsilon^\nu(k) \Big|_{k^2=m_V^2} &= 0, \\
 \Gamma_{R,\mu\nu}^{\hat{V}\hat{V}'}(k, m_V) \varepsilon^\nu(k) \Big|_{k^2=m_V^2} &= 0, \\
 \lim_{k \rightarrow m_V^2} \frac{1}{k^2 - m_V^2} \Gamma_{R,\mu\nu}^{VV}(k, m_V) \varepsilon^\nu &= -\varepsilon_\mu(k), \\
 \Gamma_R^{\hat{H}\hat{H}}(k, m_H) \Big|_{k^2=m_H^2} &= 0 \\
 \lim_{k \rightarrow m_H^2} \frac{1}{k^2 - m_H^2} \Gamma_R^{\hat{H}\hat{H}}(k, m_H) &= 1, \\
 \Gamma_R^{f\bar{f}}(k, m_f) u_f(k) \Big|_{k^2=m_f^2} &= 0 \\
 \lim_{k \rightarrow m_f^2} \frac{\not{k} + m_f}{k^2 - m_f^2} \Gamma_R^{f\bar{f}}(k, m_f) u_f(k) &= i u_f(k).
 \end{aligned} \tag{2.8}$$

With VV we indicate either AA , ZZ or WW , while with VV' we mean either AZ or ZA . Here, the renormalized 2-point functions Γ_R for background-external particles are defined in terms of the counter-

terms and the self-energies introduced in Eq. 1.30-1.31:

$$\begin{aligned}
\Gamma_{R,\mu\nu}^{\hat{V}\hat{V}}(k, m_V, m_{V'}) &= -i \left(g_{\mu\nu} - \frac{k_\mu k_\nu}{k^2} \right) \left\{ \Sigma_T^{\hat{V}\hat{V}}(k^2) + \right. \\
&\quad \left. + \left[(k^2 - m_V^2 - \delta m_V^2) Z_{VV} + (k^2 - m_{V'}^2 - \delta m_{V'}^2) \left((Z_{V'V})^{1/2} - 1 \right)^2 \right] \right\} + \\
&\quad - i \frac{k_\mu k_\nu}{k^2} \left\{ \Sigma_L^{\hat{V}\hat{V}}(k^2) + \right. \\
&\quad \left. - \left[(m_V^2 + \delta m_V^2) Z_{VV} + (m_{V'}^2 + \delta m_{V'}^2) \left((Z_{V'V})^{1/2} - 1 \right)^2 \right] \right\}, \\
\Gamma_{R,\mu\nu}^{\hat{V}\hat{V}'}(k, m_V, m_{V'}) &= -i \left(g_{\mu\nu} - \frac{k_\mu k_\nu}{k^2} \right) \left\{ \Sigma_T^{\hat{V}\hat{V}'}(k^2) + \right. \\
&\quad \left. + \left[(k^2 - m_V^2 - \delta m_V^2) Z_{VV'}^{1/2} \left((Z_{VV'})^{1/2} - 1 \right) + \right. \right. \\
&\quad \left. \left. + (k^2 - m_{V'}^2 - \delta m_{V'}^2) Z_{V'V'}^{1/2} \left((Z_{V'V})^{1/2} - 1 \right) \right] \right\} + \\
&\quad - i \frac{k_\mu k_\nu}{k^2} \left\{ \Sigma_L^{\hat{V}\hat{V}'}(k^2) + \right. \\
&\quad \left. - \left[(m_V^2 + \delta m_V^2) Z_{VV'}^{1/2} \left((Z_{VV'})^{1/2} - 1 \right) + \right. \right. \\
&\quad \left. \left. + (m_{V'}^2 + \delta m_{V'}^2) Z_{V'V'}^{1/2} \left((Z_{V'V})^{1/2} - 1 \right) \right] \right\}, \\
\Gamma_{R,\mu\nu}^{\hat{W}\hat{W}}(k, m_W) &= -i \left(g_{\mu\nu} - \frac{k_\mu k_\nu}{k^2} \right) \left\{ \Sigma_T^{\hat{W}\hat{W}}(k^2) + [(k^2 - m_W^2 - \delta m_W^2) Z_W] \right\} \\
&\quad - i \frac{k_\mu k_\nu}{k^2} \left\{ \Sigma_L^{\hat{W}\hat{W}}(k^2) - [(m_W^2 + \delta m_W^2) Z_W] \right\}, \\
\Gamma_R^{\hat{H}\hat{H}}(k, m_H) &= i \left\{ \Sigma^{\hat{H}\hat{H}}(k^2) + (k^2 - m_H^2 - \delta m_H^2) Z_H \right\}, \\
\Gamma_{R,\mu}^{\hat{A}\hat{H}}(k, m_H) &= i k_\mu \Sigma^{\hat{A}\hat{H}}(k^2), \\
\Gamma^{f\bar{f}}(k, m_f) &= i \not{k} \omega_- \left[\Sigma_L^{f\bar{f}}(k^2) + Z_L^f \right] + i \not{k} \omega_+ \left[\Sigma_R^{f\bar{f}}(k^2) + Z_R^f \right] + \\
&\quad - i(m_f + \delta m_f) \left(Z_L^f \right)^{1/2} \left(Z_R^f \right)^{1/2} + i(m_f + \delta m_f) \Sigma_S^{f\bar{f}}(k^2). \tag{2.9}
\end{aligned}$$

2.1.3 Example diagrams

The self-energies Σ introduced in the previous Section are computed starting from bubble diagrams. The complexity of these calculations grows with the perturbative order, due to the increasing number of Feynman diagrams. In this section, we provide some examples of diagrammatic topologies that contribute to the fermion and gauge boson self-energies at one-loop EW, two-loop mixed QCD-EW, and two-loop pure EW levels.

At $\mathcal{O}(\alpha)$ the self-energy topologies are fundamentally simple, consisting of either bubble or tadpole integrals. Despite their simplicity, they form the foundation of the renormalization procedure. For a generic fermion f , the one-loop self-energy is generated by the emission and reabsorption of a gauge boson, a photon, a W or a Z , or a scalar such as the Higgs. For a vector boson V , the one-loop self-energy




Figure 2.1: Examples of diagrams which contribute to the one-loop EW self-energies.




Figure 2.2: Examples of diagrams which contribute to the two-loop QCD-EW self-energies.

receives contributions from several distinct loops. For example, they contain a closed loop of quarks or leptons. While, for gauge bosons, one must also include closed loops of gauge bosons, arising from the triple gauge coupling, and Faddeev-Popov ghosts to ensure unitarity. Examples of such diagrams are shown in Fig. 2.1. At $\mathcal{O}(\alpha_s\alpha)$ the number of diagrams is limited due to the fact that gluons enter only through their coupling to quarks. At two-loop, in general, one has to consider also one-loop EW diagrams with the insertion of one-loop QCD counterterms, or vice versa. Examples of diagrams are shown in Fig. 2.2. At $\mathcal{O}(\alpha^2)$ the number of diagrams grows exponentially. At this perturbative order we encounter two-loop bubbles with up to five massive lines. Examples of such diagrams are shown in Fig. 2.3.

2.1.4 Counter-term expressions

Inserting the expression of the 2-point renormalized Green function given in Eq. 2.9 into Eq. 2.8, and collecting different powers of α , we obtain the explicit expressions for the mass counter-terms:

$$\begin{aligned}
 \delta m_W^{2(0,1)} &= \Sigma_T^{\hat{W}\hat{W}(0,1)}(m_W^2), \\
 \delta m_Z^{2(0,1)} &= \Sigma_T^{\hat{Z}\hat{Z}(0,1)}(m_Z^2), \\
 \delta m_H^{2(0,1)} &= \Sigma_T^{\hat{H}\hat{H}(0,1)}(m_H^2) \\
 \delta m_f^{(0,1)} &= \frac{m_f}{2} \left(\Sigma_L^{f\bar{f}(0,1)}(m_f^2) + \Sigma_R^{f\bar{f}(0,1)}(m_f^2) + 2\Sigma_S^{f\bar{f}(0,1)}(m_f^2) \right), \tag{2.10}
 \end{aligned}$$




Figure 2.3: Examples of diagrams which contribute to the two-loop EW self-energies.

and for the wave-functions:

$$\begin{aligned}
 \delta Z_{AA}^{(0,1)} &= -\frac{\partial \Sigma_T^{\hat{A}\hat{A}(0,1)}(k^2)}{\partial k^2} \bigg|_{k^2=0}, & \delta Z_{AZ}^{(0,1)} &= -2 \frac{\Sigma_T^{\hat{A}\hat{Z}(0,1)}(m_Z^2)}{m_Z^2}, \\
 \delta Z_{ZZ}^{(0,1)} &= -\frac{\partial \Sigma_T^{\hat{Z}\hat{Z}(0,1)}(k^2)}{\partial k^2} \bigg|_{k^2=m_Z^2}, & \delta Z_{ZA}^{(0,1)} &= 2 \frac{\Sigma_T^{\hat{Z}\hat{A}(0,1)}(0)}{m_Z^2} = 0, \\
 \delta Z_W^{(0,1)} &= -\frac{\partial \Sigma_T^{\hat{W}\hat{W}(0,1)}(k^2)}{\partial k^2} \bigg|_{k^2=m_W^2}, & \delta Z_H^{(0,1)} &= -\frac{\partial \Sigma_T^{\hat{H}\hat{H}(0,1)}(k^2)}{\partial k^2} \bigg|_{k^2=m_H^2}, \\
 \delta Z_L^{f(0,1)} &= -\Sigma_L^{f\bar{f}(0,1)}(m_f^2) - m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma_L^{f\bar{f}(0,1)}(p^2) + \Sigma_R^{f\bar{f}(0,1)}(p^2) + 2\Sigma_S^{f\bar{f}(0,1)}(p^2) \right]_{p^2=m_f^2}, \\
 \delta Z_R^{f(0,1)} &= -\Sigma_R^{f\bar{f}(0,1)}(m_f^2) - m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma_L^{f\bar{f}(0,1)}(p^2) + \Sigma_R^{f\bar{f}(0,1)}(p^2) + 2\Sigma_S^{f\bar{f}(0,1)}(p^2) \right]_{p^2=m_f^2} \quad (2.11)
 \end{aligned}$$

These expressions are well known in literature [27]. By applying the same argument at $\mathcal{O}(\alpha^2)$, we obtain the expressions for the counter-terms at the second perturbative order for the masses:

$$\begin{aligned}
 \delta m_W^{2(0,2)} &= \Sigma_T^{\hat{W}\hat{W}(0,2)}(m_W^2) - \delta m_W^{2(0,1)} \delta Z_W^{(0,1)}, \\
 \delta m_Z^{2(0,2)} &= \Sigma_T^{\hat{Z}\hat{Z}(0,2)}(m_Z^2) - \delta m_Z^{2(0,1)} \delta Z_{ZZ}^{(0,1)} + \frac{1}{4} m_Z^2 \left(\delta Z_{AZ}^{(0,1)} \right)^2, \\
 \delta m_H^{2(0,2)} &= \Sigma_T^{\hat{H}\hat{H}(0,2)}(m_H^2) - \delta m_H^{2(0,1)} \delta Z_H^{(0,1)} \\
 \delta m_f^{(0,2)} &= \frac{m_f}{2} \left[\Sigma_L^{f\bar{f}(0,2)}(m_f^2) + \Sigma_R^{f\bar{f}(0,2)}(m_f^2) + 2\Sigma_S^{f\bar{f}(0,2)}(m_f^2) \right. \\
 &\quad \left. + \frac{1}{4} \left(\delta Z_L^{f(0,1)} \right)^2 + \frac{1}{4} \left(\delta Z_R^{f(0,1)} \right)^2 + \frac{1}{2} \delta Z_L^{f(0,1)} \delta Z_R^{f(0,1)} \right] \\
 &\quad - \frac{1}{2} \delta m_f^{(0,1)} \left[\delta Z_L^{f(0,1)} + \delta Z_R^{f(0,1)} + \Sigma_S^{f\bar{f}(0,1)}(m_f^2) \right], \quad (2.12)
 \end{aligned}$$

and for the wave-functions:

$$\begin{aligned}
\delta Z_W^{(0,2)} &= - \frac{\partial \Sigma_T^{\hat{W}\hat{W}(0,2)}(k^2)}{\partial k^2} \Big|_{k^2=m_W^2}, \\
\delta Z_{ZZ}^{(0,2)} &= - \frac{\partial \Sigma_T^{\hat{Z}\hat{Z}(0,2)}(k^2)}{\partial k^2} \Big|_{k^2=m_Z^2} - \frac{1}{4} \left(\delta Z_{AZ}^{(0,1)} \right)^2, \\
\delta Z_{AZ}^{(0,2)} &= -2 \frac{\Sigma_T^{\hat{A}\hat{Z}(0,2)}(m_Z^2)}{m_Z^2} - \frac{1}{2} \delta Z_{AA}^{(0,1)} \delta Z_{AZ}^{(0,1)} \\
&\quad - \frac{1}{4} \left(\delta Z_{AZ}^{(0,1)} \right)^2 + \frac{1}{m_Z^2} \delta m_Z^{2(0,1)} \delta Z_{AZ}^{(0,1)}, \\
\delta Z_{ZA}^{(0,2)} &= 2 \frac{\Sigma_T^{\hat{Z}\hat{A}(0,2)}(0)}{m_Z^2} - \frac{1}{2} \delta Z_{ZZ}^{(0,1)} \delta Z_{ZA}^{(0,1)} \\
&\quad - \frac{1}{4} \left(\delta Z_{ZA}^{(0,1)} \right)^2 + \frac{1}{m_Z^2} \delta m_Z^{2(0,1)} \delta Z_{ZA}^{(0,1)} = 0, \\
\delta Z_{AA}^{(0,2)} &= - \frac{\partial \Sigma_T^{\hat{A}\hat{A}(0,2)}(k^2)}{\partial k^2} \Big|_{k^2=0}, \\
\delta Z_H^{(0,2)} &= - \frac{\partial \Sigma_T^{\hat{H}\hat{H}(0,2)}(k^2)}{\partial k^2} \Big|_{k^2=m_H^2}, \\
\delta Z_L^{f(0,2)} &= -\Sigma_L^{f\bar{f}(0,2)}(m_f^2) - m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma_L^{f\bar{f}(0,2)}(p^2) + \Sigma_R^{f\bar{f}(0,2)}(p^2) + 2\Sigma_S^{f\bar{f}(0,2)}(p^2) \right]_{p^2=m_f^2} \\
&\quad + m_f \delta m_f^{(0,1)} \frac{\partial \Sigma_S^{f\bar{f}(0,1)}(m_f^2)}{\partial p^2} \Big|_{p^2=m_f^2}, \\
\delta Z_R^{f(0,2)} &= -\Sigma_R^{f\bar{f}(0,2)}(m_f^2) - m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma_L^{f\bar{f}(0,2)}(p^2) + \Sigma_R^{f\bar{f}(0,2)}(p^2) + 2\Sigma_S^{f\bar{f}(0,2)}(p^2) \right]_{p^2=m_f^2} \\
&\quad + m_f \delta m_f^{(0,1)} \frac{\partial \Sigma_S^{f\bar{f}(0,1)}(m_f^2)}{\partial p^2} \Big|_{p^2=m_f^2}. \tag{2.13}
\end{aligned}$$

We observe that at $\mathcal{O}(\alpha^2)$ the counter-terms depend both on the self-energies, and the counter-terms at the previous perturbative order. We also note that the value of $\delta Z_{ZA}^{(i)}$ in the BGM is zero due to Ward identities.

Finally, we present also the expression for the counter-terms at $\mathcal{O}(\alpha_s\alpha)$, a general discussion can be found in [73].

$$\begin{aligned}
\delta m_W^{2(1,1)} &= \Sigma_T^{\hat{W}\hat{W}(1,1)}(m_W^2), \\
\delta m_Z^{2(1,1)} &= \Sigma_T^{\hat{Z}\hat{Z}(1,1)}(m_Z^2), \\
\delta m_H^{2(1,1)} &= \Sigma_T^{\hat{H}\hat{H}(1,1)}(m_H^2) \\
\delta m_f^{(1,1)} &= \frac{m_f}{2} \left[\Sigma_L^{f\bar{f}(1,1)}(m_f^2) + \Sigma_R^{f\bar{f}(1,1)}(m_f^2) + 2\Sigma_S^{f\bar{f}(1,1)}(m_f^2) \right. \\
&\quad + \frac{1}{4}\delta Z_L^{f(0,1)}\delta Z_L^{f(1,0)} + \frac{1}{4}\delta Z_R^{f(0,1)}\delta Z_R^{f(1,0)} + \frac{1}{4}\delta Z_L^{f(0,1)}\delta Z_R^{f(1,0)} + \frac{1}{4}\delta Z_L^{f(1,0)}\delta Z_R^{f(0,1)} \Big] \\
&\quad - \frac{1}{4}\delta m_f^{(0,1)} \left[\delta Z_L^{f(0,1)} + \delta Z_R^{f(0,1)} + \Sigma_S^{f\bar{f}(0,1)}(m_f^2) \right] \\
&\quad - \frac{1}{4}\delta m_f^{(1,0)} \left[\delta Z_L^{f(1,0)} + \delta Z_R^{f(1,0)} + \Sigma_S^{f\bar{f}(1,0)}(m_f^2) \right]. \tag{2.14}
\end{aligned}$$

While for the wave-functions we have:

$$\begin{aligned}
\delta Z_W^{(1,1)} &= - \frac{\partial \Sigma_T^{\hat{W}\hat{W}(1,1)}(k^2)}{\partial k^2} \Big|_{k^2=m_W^2}, \\
\delta Z_{ZZ}^{(1,1)} &= - \frac{\partial \Sigma_T^{\hat{Z}\hat{Z}(1,1)}(k^2)}{\partial k^2} \Big|_{k^2=m_Z^2}, \\
\delta Z_{AZ}^{(1,1)} &= -2 \frac{\Sigma_T^{\hat{A}\hat{Z}(1,1)}(m_Z^2)}{m_Z^2}, \\
\delta Z_{ZA}^{(1,1)} &= 0, \\
\delta Z_{AA}^{(1,1)} &= - \frac{\partial \Sigma_T^{\hat{A}\hat{A}(1,1)}(k^2)}{\partial k^2} \Big|_{k^2=0}, \\
\delta Z_H^{(1,1)} &= - \frac{\partial \Sigma_T^{\hat{H}\hat{H}(1,1)}(k^2)}{\partial k^2} \Big|_{k^2=m_H^2}, \\
\delta Z_L^{f(1,1)} &= -\Sigma_L^{f\bar{f}(1,1)}(m_f^2) - m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma_L^{f\bar{f}(1,1)}(p^2) + \Sigma_R^{f\bar{f}(1,1)}(p^2) + 2\Sigma_S^{f\bar{f}(1,1)}(p^2) \right]_{p^2=m_f^2} \\
&\quad + \frac{1}{2}m_f\delta m_f^{(0,1)} \frac{\partial \Sigma_S^{f\bar{f}(1,0)}(m_f^2)}{\partial p^2} \Big|_{p^2=m_f^2} + \frac{1}{2}m_f\delta m_f^{(1,0)} \frac{\partial \Sigma_S^{f\bar{f}(0,1)}(m_f^2)}{\partial p^2} \Big|_{p^2=m_f^2}, \\
\delta Z_R^{f(1,1)} &= -\Sigma_R^{f\bar{f}(1,1)}(m_f^2) - m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma_L^{f\bar{f}(1,1)}(p^2) + \Sigma_R^{f\bar{f}(1,1)}(p^2) + 2\Sigma_S^{f\bar{f}(1,1)}(p^2) \right]_{p^2=m_f^2} \\
&\quad + \frac{1}{2}m_f\delta m_f^{(0,1)} \frac{\partial \Sigma_S^{f\bar{f}(1,0)}(m_f^2)}{\partial p^2} \Big|_{p^2=m_f^2} + \frac{1}{2}m_f\delta m_f^{(1,0)} \frac{\partial \Sigma_S^{f\bar{f}(0,1)}(m_f^2)}{\partial p^2} \Big|_{p^2=m_f^2}. \tag{2.15}
\end{aligned}$$

2.1.5 Electric charge renormalization

The renormalization of the electric charge at the two-loop electroweak level has been discussed in detail in Ref. [79]. Specifically, the renormalized electric charge e is defined by imposing the Thomson

and a field excitation $\hat{H}(x)$ does not provide an expansion of the effective Higgs potential about its true minimum. The location of the minimum is quantified by the vev v , which itself is not a free parameter, but determined by the free parameters of the theory. Minimizing the Higgs potential to define v in perturbation theory necessarily leads to tadpole-like contributions beyond tree level, just by the perturbative ordering and truncation of contributions. Technically, it is desirable to organize the perturbative bookkeeping by appropriate parameter and field definition and renormalization in such a way that the occurrence of tadpole contributions is removed at higher order, thus reducing the number of diagrams one has to compute. In order to achieve this, we start from the renormalized 1-point vertex function $\Gamma_R^{\hat{H}}$:

$$\Gamma_R^{\hat{H}} = T^{\hat{H}} + (t + \delta t) Z_H^{1/2}, \quad \text{with} \quad t = v \left(\mu^2 - \lambda \frac{v^2}{4} \right) \quad (2.20)$$

and we impose that $\Gamma_R^{\hat{H}} = 0$ at any perturbative order. We can therefore obtain the explicit formulae for δt :

$$\begin{aligned} \delta t^{(0,1)} &= -T^{\hat{H}(0,1)}, \\ \delta t^{(0,2)} &= -T^{\hat{H}(0,2)} + \frac{1}{2} T^{\hat{H}(0,1)} \delta Z_H^{(0,1)}. \end{aligned} \quad (2.21)$$

The renormalization condition $\Gamma_R^{\hat{H}} = 0$, ensures that wherever there is a tadpole contribution of the type of Eq. 2.19, there is a corresponding counter-term δt of the type:

$$\hat{H} \text{ --- } \otimes \quad (2.22)$$

which cancels it exactly. This implies that, in practice, all tadpole diagrams can be neglected in higher order calculations. Their contribution, however, still needs to be computed since it enters in the Higgs and Goldstone-boson mass counter-term. This procedure goes under the name of Parameter-Renormalised Tadpole Scheme and it has been discussed in [27, 48, 91, 96].

2.2 Input schemes

In any quantum field theory, the bare Lagrangian is formulated in terms of fundamental parameters, such as couplings, masses, and mixing angles. However, these bare parameters are formally divergent and do not possess a direct physical meaning on their own. To make concrete theoretical predictions, one must establish a rigorous framework that maps these abstract quantities onto a set of finite, well-defined physical parameters. This specific choice of independent parameters, together with their precise theoretical definitions via corresponding renormalization conditions, is referred to as an input scheme. Crucially, an input scheme is defined independently of the specific experimental observables that are ultimately used to extract the numerical values of these parameters.

This is particularly relevant in the EW sector of the SM. If we set aside the Higgs mass parameter for a moment, the interactions of the EW gauge sector are governed entirely by three fundamental parameters in the Lagrangian: the two gauge couplings g_1 and g_2 , and the vacuum expectation value of the Higgs field, v . Phenomenologically, however, we do not measure g_1 , g_2 , and v directly. Instead, physics at the electroweak scale is described using four highly precise, physically measurable quantities: the electromagnetic fine-structure constant, α , the Fermi constant, G_F , and the masses of the gauge bosons, m_W and

m_Z ¹. Because these four physical quantities are all derived from the same three underlying Lagrangian parameters, they are not strictly independent and relations between them exist, over-constraining the system. For this reason, we cannot treat all four experimental quantities as independent inputs. We must select a set of three independent parameters to fully determine the EW sector. The remaining quantity is then treated not as an input, but as a derived, predictable observable.

In a hypothetical, all-orders calculation, the theoretical prediction for any physical process is strictly invariant regardless of which three parameters are chosen. However, in practical applications the perturbation series is truncated. This truncation breaks exact scheme invariance, meaning the numerical value of a prediction will carry a residual dependence on the chosen input parameters.

More importantly, the choice of these three parameters dictates how higher-order quantum corrections are distributed throughout the calculation. A given input scheme might effectively absorb large, universal corrections, such as those arising from the running of couplings across disparate energy scales, directly into the tree-level definitions of the parameters. This prevents these large terms from appearing explicitly in the virtual corrections, thereby guaranteeing a fast and stable convergence of the perturbation series. Conversely, an ill-suited scheme choice can lead to artificially inflated higher-order corrections, rendering the theoretical prediction slow to converge or unreliable.

In principle, various combinations of precision observables can be selected to construct a valid input scheme. In this thesis, we focus specifically on schemes that adopt the gauge boson masses m_W and m_Z as two of the primary inputs. Within this class of schemes, the choice of the third parameter defines the specific input scheme employed.

2.2.1 $\alpha(0)$ scheme

The most intuitive choice is the $\alpha(0)$ scheme, which uses the fine-structure constant in the Thomson limit as the third input parameter. The independent parameters are therefore $\{\alpha(0), m_Z, m_W\}$. In this scheme, the weak mixing angle is defined to all orders by the ratio of the gauge boson masses $\sin^2 \theta_W = m_W^2/m_Z^2$. The corresponding charge renormalization constant, δZ_e , is fixed by requiring that the $ee\gamma$ vertex reduces to $ie\gamma^\mu$ in the Thomson limit, i.e. $q \rightarrow 0$. While this scheme is exact and physically transparent for processes involving real external photons, it suffers from severe drawbacks for high-energy weak processes. The running of the electromagnetic coupling from $q^2 = 0$ to the electroweak scale $q^2 = m_Z^2, m_W^2$ introduces large logarithmic corrections. In the $\alpha(0)$ scheme, these large logarithms contribute to the higher-order virtual corrections, leading to a slowly converging perturbative expansion.

2.2.2 G_μ scheme

To avoid the large logarithms associated with the light fermion masses, it is highly advantageous to choose an input parameter that is naturally defined at the electroweak scale. The G_μ scheme accomplishes this by utilizing the Fermi constant, G_F , which is precisely extracted from the muon decay width Γ_μ . The input parameters are therefore $\{G_F, m_W, m_Z\}$. Rather than being a fundamental SM parameter, G_F is defined within the low-energy Fermi effective theory as the coupling strength of the four-fermion contact interaction. By requiring that the full loop-corrected SM amplitude for the muon decay matches with the Fermi theory prediction, one establishes a relation between the fine-structure constant and G_F :

$$\alpha_{G_\mu} = \frac{\sqrt{2}G_F m_W^2 \sin^2 \theta_W}{\pi} (1 - \Delta r). \quad (2.23)$$

¹In addition to these four quantities, the effective weak mixing angle $\sin^2 \theta_W$, represents another critical precision observable. It plays a particularly important role as a phenomenological input for characterizing Neutral-Current Drell-Yan processes.

Here, Δr encapsulates the radiative corrections to muon decay, including the large logarithms from the running of α and the dominant quadratic top-quark mass dependence. By expressing the leading-order amplitudes in terms of α_{G_μ} rather than $\alpha(0)$, the G_μ scheme effectively absorbs these large, universal electroweak corrections into the tree-level couplings. Consequently, the explicit higher-order EW corrections evaluated in Chapters 6 will be significantly smaller and more perturbatively stable when computing Drell-Yan observables.

2.3 Complex Mass Scheme

In perturbation theory, unstable particles such as the W and Z bosons manifest as poles in Feynman propagators. When the invariant mass squared p^2 of a propagating particle approaches its real mass squared m^2 , the standard tree-level propagator $1/(p^2 - m^2)$ diverges. To evaluate matrix elements and cross-sections near such resonances, it is necessary to regularize the pole by resumming the self-energies. This Dyson resummation shifts the pole into the complex plane by introducing a decay width Γ , thus modifying the denominator to $(p^2 - m^2 + im\Gamma)$. While mathematically regulating the pole, this straightforward resummation mixes different orders of perturbation theory. In non-abelian gauge theories like the SM, gauge invariance relies on exact, order-by-order cancellations governed by Slavnov-Taylor and Ward identities. The partial inclusion of higher-order terms via the width breaks these delicate cancellations, violating gauge invariance. This violation can manifest as highly gauge-dependent predictions or the breakdown of unitarity at high scattering energies. The Complex Mass Scheme (CMS) [97–99] was introduced to provide a universal, gauge-invariant framework to treat resonances in higher-order calculations. In particular, the mass of a gauge bosons is considered as complex quantity, and is defined as the location of the pole of the corresponding propagator. The latter, is a gauge invariant quantity:

$$\mu^2 = m^2 - im\Gamma. \quad (2.24)$$

By treating μ^2 as the fundamental mass parameter, the algebraic structure of the SM Lagrangian remains strictly intact. Because all underlying symmetries of the theory are preserved under the analytic continuation of parameters into the complex plane, the CMS exactly preserves all Slavnov-Taylor and Ward identities at any given order in perturbation theory. Consequently, the scheme avoids unitarity violations and spurious gauge dependences in resonant and non-resonant phase-space regions simultaneously.

2.3.1 Impact on renormalization

The transition to the CMS requires extending the traditional on-shell renormalization procedure into the complex plane. In this framework, the bare mass squared is split into a complex renormalized mass squared and a corresponding complex counterterm, $\mu_{V,0}^2 = \mu_V^2 + \delta\mu_V^2$. Unlike the standard on-shell scheme, where renormalization conditions force the renormalized mass to equal the physical mass by taking only the real part of the self-energy, the CMS fixes the complex mass exactly to the location of the complex pole of the propagator. Consequently, the mass counterterm evaluates the unrenormalized self-energy strictly at this complex pole without projecting out the real part, meaning the traditional real-part projection operation is entirely omitted. Furthermore, the wave-function renormalization constants for unstable fields also become complex quantities to properly normalize the residues of these complex poles. In the CMS, hence, we need to evaluate the unrenormalized self-energies $\Sigma(p^2)$ at the complex pole $p^2 = \mu^2$. In order to do so, we directly have to compute the loop integrals for complex values of the squared momentum. This requires the analytic continuation of the standard scalar integrals that appear in

the self-energies into the complex plane. Because both the internal masses of the virtual particles and the external invariant mass are complex, the integration must carefully track the branch cuts and Riemann sheets of the logarithms and dilogarithms involved. This can be done, without any approximation, for example using the series expansion approach, together with the techniques that will be introduced in Section 3.5 for performing the analytic continuation of the result into the complex plane.

2.3.2 Impact on input schemes

The adoption of complex masses inherently propagates to derived parameters and the definition of electroweak input schemes. For instance, the weak mixing angle is defined via the exact ratio of the complex gauge boson masses, $c_W^2 = \mu_W^2 / \mu_Z^2$, which means both c_W^2 and the derived sine $s_W^2 = 1 - c_W^2$ become complex quantities. Consequently, because couplings are algebraically linked to these masses and the mixing angle, derived quantities in the CMS inevitably acquire imaginary parts. When employing input parameter frameworks such as the G_μ scheme, the relations connecting high-precision inputs like the Fermi constant G_F , μ_Z and μ_W to the tree-level couplings must consistently utilize these complex values. This consistent adoption of complex parameters ensures that the exact algebraic relations necessary for gauge cancellations are maintained throughout the entire calculation.

2.4 γ^5 problem

One of the biggest issues when computing higher order amplitudes is the so-called γ^5 problem. This originates from the fact that the computation of multi-loop Feynman integrals is usually performed in dimensional regularization (DR) (see Chapter 3), i.e. in $d = 4 - 2\epsilon$ dimensions. The latter, however, must be adopted also in the evaluation of Lorentz and Dirac algebra.

Some fundamental relations can be easily generalized to d dimensions. For example we find:

$$\begin{aligned} g^{\mu\nu} g_{\mu\nu} &= d, \\ \{\gamma_\mu, \gamma_\nu\} &= 2g_{\mu\nu}, \\ \gamma_\mu \gamma^\mu &= d. \end{aligned} \tag{2.25}$$

On the other hand, when trying to generalize the definitions involving $\gamma_5 = -i\gamma_0\gamma_1\gamma_2\gamma_3$ to d dimensions one finds some ambiguities. This is because while DR preserves Lorentz and gauge invariance, it does not preserve chiral invariance. This is due to the fact that γ_5 is inherently a 4-dimensional object, and there is no canonical way to continue its 4 dimensional definition to d -dimensions. On an algebraic level, it is impossible to define a d -dimensional algebra that obeys both the anticommutation relation

$$\{\gamma_\mu, \gamma_5\} = 0 \tag{2.26}$$

and the trace condition

$$\text{Tr}(\gamma^\mu \gamma^\nu \gamma^\rho \gamma^\sigma \gamma^5) = -4i\epsilon^{\mu\nu\rho\sigma}, \tag{2.27}$$

where $\epsilon^{\mu\nu\rho\sigma}$ is the Levi-Civita totally anti-symmetric tensor. This can be seen explicitly by considering the trace in Eq. 2.27 and inserting the identity $\gamma_\alpha \gamma^\alpha = d$:

$$\frac{1}{d} \text{Tr}(\gamma^\mu \gamma^\nu \gamma^\rho \gamma^\sigma \gamma^5 \gamma_\alpha \gamma^\alpha) = \frac{1}{d} \text{Tr}(\gamma^\alpha \gamma^\mu \gamma^\nu \gamma^\rho \gamma^\sigma \gamma^5 \gamma_\alpha), \tag{2.28}$$

where we used the cyclicity of the trace to bring the γ^α in front. Now by using the standard algebra relations $\{\gamma_\mu, \gamma_\nu\} = 2g_{\mu\nu}$ and $\{\gamma_\mu, \gamma_5\} = 0$ we arrive at:

$$(d - 4)\text{Tr}(\gamma^\mu \gamma^\nu \gamma^\rho \gamma^\sigma \gamma^5) = 0, \quad (2.29)$$

which implies that the trace is 0, in contradiction with Eq. 2.27.

Since γ_5 and $\varepsilon_{\mu\nu\rho\sigma}$ are intrinsically four-dimensional objects, it is convenient to identify a four-dimensional subspace and sacrifice the construction of a manifestly Lorentz-covariant formalism. Following this spirit, we write

$$k^\mu = \bar{k}^\mu + \hat{k}^\mu, \quad \gamma^\mu = \bar{\gamma}^\mu + \hat{\gamma}^\mu, \quad (2.30)$$

where bars denote four-dimensional objects and hats denote (-2ϵ) -dimensional objects. Similarly, we write the usual four-dimensional identities,

$$\varepsilon^{\mu\nu\rho\sigma} \varepsilon_{\mu\nu\rho\sigma} = -24, \quad (2.31)$$

$$\varepsilon^{\mu\nu\rho}{}_\alpha \varepsilon_{\mu\nu\rho\beta} = -6 \bar{g}_{\alpha\beta}, \quad (2.32)$$

$$\varepsilon^{\mu\nu}{}_{\alpha\delta} \varepsilon_{\mu\nu\beta\eta} = -2 (\bar{g}_{\alpha\beta} \bar{g}_{\delta\eta} - \bar{g}_{\alpha\eta} \bar{g}_{\beta\delta}), \quad (2.33)$$

$$\begin{aligned} \varepsilon^\mu{}_{\alpha\delta\kappa} \varepsilon_{\mu\beta\eta\lambda} = & \bar{g}_{\alpha\lambda} \bar{g}_{\beta\kappa} \bar{g}_{\delta\eta} - \bar{g}_{\alpha\lambda} \bar{g}_{\beta\delta} \bar{g}_{\eta\kappa} - \bar{g}_{\alpha\eta} \bar{g}_{\beta\kappa} \bar{g}_{\delta\lambda} \\ & + \bar{g}_{\alpha\beta} \bar{g}_{\delta\lambda} \bar{g}_{\eta\kappa} + \bar{g}_{\alpha\eta} \bar{g}_{\beta\delta} \bar{g}_{\eta\kappa} - \bar{g}_{\alpha\beta} \bar{g}_{\delta\eta} \bar{g}_{\eta\kappa}, \end{aligned} \quad (2.34)$$

$$\begin{aligned} \varepsilon_{\alpha\delta\kappa\tau} \varepsilon_{\beta\eta\lambda\chi} = & -\bar{g}_{\alpha\chi} \bar{g}_{\beta\tau} \bar{g}_{\delta\lambda} \bar{g}_{\eta\kappa} + \bar{g}_{\alpha\chi} \bar{g}_{\beta\tau} \bar{g}_{\delta\eta} \bar{g}_{\kappa\lambda} + \bar{g}_{\alpha\chi} \bar{g}_{\beta\kappa} \bar{g}_{\delta\lambda} \bar{g}_{\eta\tau} - \bar{g}_{\alpha\chi} \bar{g}_{\beta\delta} \bar{g}_{\eta\tau} \bar{g}_{\kappa\lambda} \\ & -\bar{g}_{\alpha\lambda} \bar{g}_{\beta\kappa} \bar{g}_{\delta\eta} \bar{g}_{\eta\tau} + \bar{g}_{\alpha\lambda} \bar{g}_{\beta\delta} \bar{g}_{\eta\tau} \bar{g}_{\kappa\lambda} + \bar{g}_{\alpha\lambda} \bar{g}_{\beta\tau} \bar{g}_{\delta\chi} \bar{g}_{\eta\kappa} - \bar{g}_{\alpha\eta} \bar{g}_{\beta\tau} \bar{g}_{\delta\chi} \bar{g}_{\eta\kappa} \\ & -\bar{g}_{\alpha\lambda} \bar{g}_{\beta\kappa} \bar{g}_{\delta\chi} \bar{g}_{\eta\tau} + \bar{g}_{\alpha\beta} \bar{g}_{\delta\chi} \bar{g}_{\eta\tau} \bar{g}_{\kappa\lambda} + \bar{g}_{\alpha\eta} \bar{g}_{\beta\kappa} \bar{g}_{\delta\chi} \bar{g}_{\eta\tau} - \bar{g}_{\alpha\beta} \bar{g}_{\delta\chi} \bar{g}_{\eta\kappa} \bar{g}_{\eta\tau} \\ & -\bar{g}_{\alpha\lambda} \bar{g}_{\beta\tau} \bar{g}_{\delta\eta} \bar{g}_{\eta\kappa} + \bar{g}_{\alpha\eta} \bar{g}_{\beta\tau} \bar{g}_{\delta\lambda} \bar{g}_{\eta\kappa} + \bar{g}_{\alpha\lambda} \bar{g}_{\beta\delta} \bar{g}_{\eta\tau} \bar{g}_{\kappa\chi} - \bar{g}_{\alpha\beta} \bar{g}_{\delta\lambda} \bar{g}_{\eta\tau} \bar{g}_{\kappa\chi} \\ & -\bar{g}_{\alpha\eta} \bar{g}_{\beta\delta} \bar{g}_{\eta\kappa} \bar{g}_{\eta\tau} + \bar{g}_{\alpha\beta} \bar{g}_{\delta\eta} \bar{g}_{\eta\kappa} \bar{g}_{\eta\tau} + \bar{g}_{\alpha\lambda} \bar{g}_{\beta\kappa} \bar{g}_{\delta\eta} \bar{g}_{\eta\tau} - \bar{g}_{\alpha\lambda} \bar{g}_{\beta\delta} \bar{g}_{\eta\kappa} \bar{g}_{\eta\tau} \\ & -\bar{g}_{\alpha\eta} \bar{g}_{\beta\kappa} \bar{g}_{\delta\lambda} \bar{g}_{\eta\tau} + \bar{g}_{\alpha\beta} \bar{g}_{\delta\lambda} \bar{g}_{\eta\kappa} \bar{g}_{\eta\tau} + \bar{g}_{\alpha\eta} \bar{g}_{\beta\delta} \bar{g}_{\eta\kappa} \bar{g}_{\eta\tau} - \bar{g}_{\alpha\beta} \bar{g}_{\delta\eta} \bar{g}_{\eta\kappa} \bar{g}_{\eta\tau}, \end{aligned} \quad (2.35)$$

using $\bar{g}_{\mu\nu}$ to denote the four-dimensional metric tensor. In the following, we will focus on two different prescriptions for treating γ_5 in d dimensions.

2.4.1 HVBM's prescription

The first scheme presented is the one proposed by 't Hooft-Veltman-Breitenlohner-Maison [80–83]. In this scheme the chiral couplings are treated in a consistent way in d dimensions by retaining the four-dimensional definition of γ_5 , Eq. 2.27, at the expense of its anti-commutation relation with the standard Dirac matrices. In place of Eq. 2.26, one can write:

$$\{\bar{\gamma}_\mu, \gamma_5\} = 0, \quad [\hat{\gamma}_\mu, \gamma_5] = 0, \quad \{\bar{\gamma}_\mu, \hat{\gamma}_\nu\} = 0. \quad (2.36)$$

HVBM's scheme allows for a very familiar treatment of the usual families of Dirac traces. The trace of an even number of Dirac matrices with d dimensional indices may be conveniently evaluated using the recursive formula

$$\text{Tr}(\gamma_{\nu_1} \cdots \gamma_{\nu_n}) = \sum_{i=1}^{n-1} (-1)^{i+1} g_{\nu_1 \nu_{i+1}} \text{Tr}(\gamma_{\nu_2} \cdots \gamma_{\nu_i} \gamma_{\nu_{i+2}} \cdots \gamma_{\nu_n}), \quad (2.37)$$

valid for even n , with termination criterion

$$\text{Tr}(\mathbf{1}) = 4. \quad (2.38)$$

Eq. 2.37 is derived by repeatedly applying Eq. 2.26 to move the γ_{ν_1} factor all the way to the right. Finally, the cyclic property of the trace is required to identify the final term, $-\text{Tr}(\gamma_{\nu_2} \cdots \gamma_{\nu_n} \gamma_{\nu_1})$, with the original expression on the left-hand side.

The trace

$$\text{Tr}(\gamma_{\nu_1} \cdots \gamma_{\nu_n} \gamma_5) \quad (2.39)$$

for even n may be evaluated in much the same way by replacing γ_5 by the covariant version of its definition,

$$\gamma_5 = -\frac{i}{4!} \varepsilon^{\mu\nu\rho\sigma} \gamma_\mu \gamma_\nu \gamma_\rho \gamma_\sigma, \quad (2.40)$$

and then evaluating the resulting trace of $n + 4$ Dirac matrices using Eqs. 2.37 and 2.38.

HVBM's prescription for γ_5 has the very unfortunate drawback that the SM's chiral symmetry is violated in all bare perturbative calculations involving chiral couplings. As was originally pointed out by Breitenlohner and Maison in [80–82] and reviewed in [100], it is possible to restore the chiral symmetry of the Standard Model as part of the renormalization program by introducing various finite counterterms on top of the usual divergent ones. Let us emphasize here that, in general, one expects all interactions in the SM to require finite counterterms at some order in perturbation theory, not just those which explicitly involve γ_5 . Furthermore, these finite counterterms must generically be calculated beyond $\mathcal{O}(\epsilon^0)$ to consistently restore chiral symmetry at higher orders.

2.4.2 Kreimer's prescription

In the γ_5 scheme proposed by Kreimer, chiral couplings are handled consistently in d dimensions by preserving the strictly four-dimensional definition, $\gamma_5 = -i\gamma_0\gamma_1\gamma_2\gamma_3$, though this requires abandoning the cyclicity of the Dirac trace. As argued in [84], retaining trace cyclicity is unnatural in d dimensions because the Dirac algebra inevitably becomes infinite-dimensional. Consequently, Kreimer's approach replaces the conventional linear algebra trace with specialized, non-cyclic trace operations defined by a reading point prescription, i.e. one must systematically start writing the trace from a fixed external momentum. In practice, this means that for all diagrams contributing to a given physical process, the fermion lines must be "read" starting from the exact same designated external leg.

Compared to the HVBM scheme, Kreimer's approach offers notable practical benefits. Most importantly, evaluating loop amplitudes does not alter the algebraic structure of any SM Feynman rule. By adopting a consistent reading point prescription, chiral symmetry is naturally maintained throughout bare perturbative evaluations, completely eliminating the need for finite counterterms and greatly simplifying the renormalization process. Furthermore, while computing box-type diagrams in Kreimer's scheme ultimately requires splitting the loop momenta as in Eq. 2.30, the majority of the numerator algebra can be processed before this split is introduced. This is a significant improvement over the HVBM scheme, which necessitates the immediate introduction of split loop momenta to efficiently facilitate algebraic simplifications.

As alluded to above, the standard families of odd Dirac traces still vanish identically,

$$\text{Tr}(\gamma_{\nu_1} \cdots \gamma_{\nu_{2n+1}}) = 0 \quad (2.41)$$

$$\text{Tr}(\gamma_{\nu_1} \cdots \gamma_{\nu_{2n+1}} \gamma_5) = 0. \quad (2.42)$$

For the standard families of even Dirac traces, [101] defines

$$\text{Tr}(\gamma_{\nu_1} \cdots \gamma_{\nu_{2n}}) = 4 \sum_{\sigma} (-1)^{\text{sgn}(\sigma)} g_{\nu_{i_1} \nu_{j_1}} \cdots g_{\nu_{i_n} \nu_{j_n}} \quad (2.43)$$

$$\begin{aligned} \text{Tr}(\gamma_{\nu_1} \cdots \gamma_{\nu_{2n}} \gamma_5) = & -\frac{i}{3!} \varepsilon_{\nu_{2n+1} \nu_{2n+2} \nu_{2n+3} \nu_{2n+4}} \\ & \sum_{\sigma} (-1)^{\text{sgn}(\sigma)} g_{\nu_{i_1} \nu_{j_1}} \cdots g_{\nu_{i_{n+2}} \nu_{j_{n+2}}}, \end{aligned} \quad (2.44)$$

where σ is the set of permutations of the $2m$ indices (with $m = n$ or $m = n + 2$) which appear in the sums on the right-hand sides, subject to the restrictions

$$1 = i_1 < \dots < i_m, \quad i_k < j_k, \quad \text{for all } k \in \{1 \dots m\}. \quad (2.45)$$

It is worth noting that these definitions are entirely equivalent to what one would calculate in the HVBM γ_5 scheme by recursively applying Tr along with Eqs. 2.37, 2.38, and 2.40. This equivalence is expected: because a chain of standard Dirac matrices carries no explicit dependence on γ_5 , Eqs. 2.37 and 2.38 must inherently reduce to the same result as Eq. 2.43.

2.4.3 d -dimensional Feynman integrals with four-dimensional scalar numerators

In this section, we outline our strategy for managing the difficult numerator terms that frequently appear in d -dimensional Feynman integrals within Standard Model perturbation theory—specifically, the four-dimensional contractions of d -dimensional loop momenta, $\bar{g}_{\mu\nu} q_i^\mu q_j^\nu$. As previously stated, Eq. 2.30, we separate the loop momenta into their respective four-dimensional and (-2ϵ) -dimensional parts, such that $q_i = \bar{q}_i + \hat{q}_i$. Following the method pioneered by Bern and Morgan [102], this decomposition effectively translates those problematic four-dimensional numerators into standard inverse propagators alongside $\hat{q}_i \cdot \hat{q}_j$ insertions. In the remainder of this section, we detail how to evaluate these insertions at the one- and two-loop levels by recasting them into dimensionally shifted integrals. For the two-loop case specifically, we review the original algorithm introduced in Ref. [103]. Alternative strategies for handling these two-loop insertions are available in the literature [104].

The one-loop Feynman integrals we are focusing on are of the form

$$\mathcal{I}_n^{(1)}[(\hat{q}_1 \cdot \hat{q}_1)^r] \equiv \int \frac{d^d q_1}{i\pi^{d/2}} \frac{(\hat{q}_1 \cdot \hat{q}_1)^r}{D_1^{\nu_1} \cdots D_n^{\nu_n}}, \quad (2.46)$$

where the D_i are the independent propagators of the integral family. Crucially, Eq. 2.46 can be rewritten in terms of standard dimensionally-shifted one-loop Feynman integrals with the same propagator structure. By absorbing the $(\hat{q}_1 \cdot \hat{q}_1)^r$ numerator insertion into the integration measure, it is shown in [102] that

$$\mathcal{I}_n^{(1)}[(\hat{q}_1 \cdot \hat{q}_1)^r] = (-1)^r \frac{\Gamma(r - \epsilon)}{\Gamma(-\epsilon)} \int \frac{d^{d+2r} q_1}{i\pi^{(d+2r)/2}} \frac{1}{D_1^{\nu_1} \cdots D_n^{\nu_n}}. \quad (2.47)$$

Even though Eq. 2.47 fundamentally relies on dimensionally-shifted integrals, its structure is completely independent of the number of external legs, the kinematic complexity, or the specific values of ν_i . For an integral topology containing t propagator denominators, the first Symanzik polynomial takes on a universally simple form at one loop:

$$\mathcal{U}_t^{(1)} = \alpha_1 + \cdots + \alpha_t, \quad (2.48)$$

a result that follows directly from the unique and trivial nature of the one-loop vacuum (tadpole) integral topology.

The most straightforward way to express the general Schwinger representation requires first sorting the Feynman propagators so that their exponents, in the $\epsilon \rightarrow 0$ limit, create a monotonically decreasing sequence. Under this limit, we assume the first n^+ propagators have positive exponents, while the subsequent n^- propagators have negative exponents. Without losing generality, we can treat the exponents of these final n^- propagators as non-positive integers. Using the Symanzik polynomials for the integral family, $\overline{\mathcal{U}}$ and $\overline{\mathcal{F}}$, the representation is given by:

$$\begin{aligned} \mathcal{I}_n^{(L)} = & e^{-i\pi \sum_{j=1}^{n^+} \nu_j} \prod_{i=1}^{n^+} \left(\int_0^\infty d\alpha_i \frac{\alpha_i^{\nu_i-1}}{\Gamma(\nu_i)} \right) \times \\ & \times \prod_{\ell=n^++1}^n \frac{\partial^{|\nu_\ell|}}{\partial \alpha_\ell^{|\nu_\ell|}} \left(\overline{\mathcal{U}}^{-d_0/2+\epsilon} e^{-\overline{\mathcal{F}}/\overline{\mathcal{U}}} \right) \Big|_{\alpha_{n^++1}=\cdots=\alpha_{n^++n^-}=0} \end{aligned} \quad (2.49)$$

This holds for $d = d_0 - 2\epsilon$, assuming d_0 is an even positive integer, with $d_0 = 4$ being our primary focus. If there are no inverse propagators present, this formula naturally reduces to the standard result expressed via the Symanzik polynomials of the specific integral topology, $\mathcal{U}$ and $\mathcal{F}$.

As highlighted in [102], a generic d -dimensional Feynman propagator at one loop takes the form $\bar{q}_1 \cdot \bar{q}_1 + 2P_{ext} \cdot \bar{q}_1 + P_{ext} \cdot P_{ext} - m^2 + \hat{q}_1 \cdot \hat{q}_1$, where P_{ext} represents a sum of four-dimensional external momenta and m is the particle mass. This specific structure justifies splitting the loop momentum integral into distinct four-dimensional and (-2ϵ) -dimensional components when deriving the standard Schwinger parametrization. Because these insertions are inherently designed to yield Feynman integrals suppressed by explicit powers of ϵ , the authors of [103] concluded that the (-2ϵ) -dimensional portion of the momentum integral must be proportional to the ϵ -dependent factor in Eq. 2.49:

$$\int d^{-2\epsilon} \hat{q}_1 e^{\overline{\mathcal{U}}_n^{(1)} \hat{q}_1 \cdot \hat{q}_1} \propto \left(\overline{\mathcal{U}}_n^{(1)} \right)^\epsilon, \quad (2.50)$$

Consequently, this term acts as a generating function for all these insertions through repeated differentiation with respect to $\overline{\mathcal{U}}_n^{(1)}$. Taking the r -th derivative of the right-hand side of this proportionality with respect to $\overline{\mathcal{U}}_n^{(1)}$ yields:

$$(-1)^r \frac{\Gamma(r-\epsilon)}{\Gamma(-\epsilon)} \left(\overline{\mathcal{U}}_n^{(1)} \right)^{\epsilon-r}, \quad (2.51)$$

By noting that the transformation $\epsilon \rightarrow \epsilon - r$ is strictly equivalent to the dimension shift $d_0 \rightarrow d_0 + 2r$ in Eq. 2.49, we immediately recover the anticipated form of Eq. 2.47.

Fortunately, the methodology outlined above extends effectively to the evaluation of generic integrals at the two-loop level. This is largely due to the existence of a uniquely defined non-factorizable two-loop tadpole topology (see Figure 2.4). A key distinction at the two-loop level is that dimension shifts are now accompanied by shifts in the propagator exponents. Here, generic non-zero insertions are formed




Figure 2.4: The unique (up to internal mass assignment) non-factorizable (left) and factorizable (right) two-loop vacuum integral topologies.

by multiplying the components $\hat{q}_1 \cdot \hat{q}_1$, $\hat{q}_2 \cdot \hat{q}_2$, and $(\hat{q}_1 - \hat{q}_2)^2$. Because the non-factorizable two-loop tadpole topology is unique, the first Symanzik polynomial for any arbitrary non-factorizable two-loop integral family takes the following specific form:

$$\overline{\mathcal{U}}_n^{(2)} = T_{q_1^2} T_{q_1 \cdot q_2} + T_{q_2^2} T_{q_1 \cdot q_2} + T_{q_1^2} T_{q_2^2}, \quad (2.52)$$

In this expression, the T variables represent sums of Schwinger parameters associated with specific propagators: $T_{q_1 \cdot q_2}$ are propagators depending on the dot product $q_1 \cdot q_2$, $T_{q_i^2}$ are propagators depending on q_i^2 but independent of $q_1 \cdot q_2$. Consequently, the two-loop equivalent of Eq. 2.50 emerges as a generating function:

$$\int d^{-2\epsilon} \hat{q}_1 d^{-2\epsilon} \hat{q}_2 e^{T_{q_1^2} \hat{q}_1 \cdot \hat{q}_1 + T_{q_2^2} \hat{q}_2 \cdot \hat{q}_2 + T_{q_1 \cdot q_2} (\hat{q}_1 - \hat{q}_2)^2} \propto \left(\overline{\mathcal{U}}_n^{(2)} \right)^\epsilon, \quad (2.53)$$

By differentiating this function with respect to $T_{q_1^2}$, $T_{q_2^2}$, and $T_{q_1 \cdot q_2}$, one can generate the necessary terms. Just as in the one-loop case, comparing the differentiated terms of Eq. 2.53 against the structure of Eq. 2.49 allows us to rewrite any arbitrary two-loop integral as a linear combination of standard Feynman integrals.

Beginning at the two-loop level, our preference for expressing numerator polynomials as linear combinations of inverse propagators and particle masses introduces a subtle complication. Specifically, if a Feynman integral's numerator contains both an explicit $\hat{q}_1 \cdot \hat{q}_1$ insertion and a propagator dependent solely on q_1 , the previously established logic for eliminating the $\hat{q}_1 \cdot \hat{q}_1$ factor breaks down. This occurs because the q_1^2 term within the propagator intrinsically includes $\hat{q}_1 \cdot \hat{q}_1$, as dictated by the decomposition $q_1^2 = \bar{q}_1 \cdot \bar{q}_1 + \hat{q}_1 \cdot \hat{q}_1$. To navigate this in practice, we found it highly effective to temporarily postpone the conversion of four-dimensional scalar products, $q_i \cdot p_j$, into inverse propagators. Since these $q_i \cdot p_j$ scalar products exist strictly in four dimensions, we can safely apply Eq. 2.53 to clear out all explicit terms before finally rewriting the leftover scalar products as linear combinations of propagators.

To clarify the procedure described above, we consider a non-trivial two-loop integral which may be directly rewritten using the technology of [103]:

$$\begin{aligned} & \text{Diagram: A two-loop integral with two external lines meeting at a vertex, forming a loop with a vertical line segment.} \\ & [\mathcal{D}_7 (\hat{q}_1 \cdot \hat{q}_1)^2] \\ &= \int \frac{d^d q_1 d^d q_2}{(i\pi^{d/2})^2} \frac{(q_2 - p_1 - p_2)^2 (\hat{q}_1 \cdot \hat{q}_1)^2}{q_1^2 q_2^2 (q_1 - q_2)^2 (q_1 - p_1)^2 (q_1 - p_1 - p_2)^2}, \end{aligned} \quad (2.54)$$

where we have introduced the notation $\mathcal{D}_7 \equiv (q_2 - p_1 - p_2)^2$ on the left-hand side for later convenience. Note that, in this example, there is no clash between the explicit factor of $(\hat{q}_1 \cdot \hat{q}_1)^2$ in Eq. 2.54 and the implicit factor of $\hat{q}_2 \cdot \hat{q}_2$ present in

$$\mathcal{D}_7 = \bar{q}_2 \cdot \bar{q}_2 - 2(p_1 + p_2) \cdot \bar{q}_2 + s + \hat{q}_2 \cdot \hat{q}_2. \quad (2.55)$$

The Schwinger parameter subsets of interest are

$$T_{q_1^2} = \alpha_1 + \alpha_4 + \alpha_6, \quad T_{q_2^2} = \alpha_2 + \alpha_7, \quad \text{and} \quad T_{q_1 \cdot q_2} = \alpha_3, \quad (2.56)$$

and differentiating the right-hand side of Eq. 2.53 twice with respect to $T_{q_1^2}$ gives $\epsilon(\epsilon - 1)$ times an insertion of

$$\frac{(T_{q_2^2} + T_{q_1 \cdot q_2})^2}{(\overline{\mathcal{U}}_6^{(2)})^2} = \frac{\alpha_2^2 + \alpha_3^2 + \alpha_7^2 + 2\alpha_2\alpha_3 + 2\alpha_2\alpha_7 + 2\alpha_3\alpha_7}{(\overline{\mathcal{U}}_6^{(2)})^2} \quad (2.57)$$

into Eq. 2.49 specialized to our two-loop example.

The presence of the denominator factor is responsible for driving a dimension shift from $4 - 2\epsilon \rightarrow 8 - 2\epsilon$. Any powers of the Schwinger parameters tied to the propagator denominators will cause positive shifts in their respective propagator exponents; specifically, raising parameter i to the power of h introduces an overall multiplicative factor of $(-1)^h \Gamma(\nu_i + h) / \Gamma(\nu_i)$. Conversely, while powers of the Schwinger parameter linked to the numerator also yield positive exponent shifts, the application of the derivatives in Eq. 2.49 produces an overall factor of $h!$ instead. Finally, because α_7 functions purely as an auxiliary parameter that is ultimately set to zero, inserting α_7^2 produces a null contribution in the particular example under consideration. Consequently, we obtain:

$$\begin{aligned} & \left[\mathcal{D}_7 (\hat{q}_1 \cdot \hat{q}_1)^2 \right] = 2\epsilon(\epsilon - 1) \left[\begin{aligned} & \left[\mathcal{D}_7 \right] \\ & + \left[\mathcal{D}_7 \right] \\ & - \left[\mathcal{D}_7 \right] \end{aligned} \right] \quad (2.58) \end{aligned}$$

The integrals on the right-hand side of Eq. 2.58 can be straightforwardly evaluated in Feynman parameters to all orders in ϵ by performing the two loop-momentum integrations analytically, one at a time, using standard one-loop formulas. After some simplifications, we obtain

$$\begin{aligned} & \left[\mathcal{D}_7 (\hat{q}_1 \cdot \hat{q}_1)^2 \right] \\ & = \frac{(5 - 2\epsilon)\Gamma^2(\epsilon)\Gamma^2(1 - \epsilon)\Gamma(2\epsilon)\Gamma(1 - 2\epsilon)\Gamma(2 - \epsilon)}{\Gamma(6 - 3\epsilon)\Gamma(3 - 2\epsilon)\Gamma^2(-1 + \epsilon)}. \quad (2.59) \end{aligned}$$

The key point of the discussion is that this procedure allows us to systematically map integrals with insertions of (-2ϵ) -dimensional scalar products, $\hat{q}_i \cdot \hat{q}_j$, onto dimensionally shifted ones. Once recast in this form, these dimensionally shifted integrals can be straightforwardly evaluated utilizing the standard and well-established techniques developed for ordinary multi-loop Feynman integrals, which will be the focus of Ch. 3. Specifically, we will employ integration-by-parts reduction algorithms to reduce them to a finite basis of master integrals. These master integrals can subsequently be evaluated via conventional analytic or numerical or semi-numerical approaches, for example using the method of differential equa-

tions. Using this pipeline we can compute a generic higher-loop amplitude consistently handling γ_5 in d dimensions.

Feynman integrals

“In climbing you are always faced with new problems in which you must perform using intuitive movements, and then later analyze them to figure out why they work.”

Wolfgang Güllich, *German pioneer of extreme free climbing.*

One of the main bottleneck for the inclusion of higher order corrections is the calculation of the so-called Feynman integrals [22, 105]. Indeed, in state-of-the-art problems their number can grow up to thousands or hundred of thousands, and hence a general and algorithmic approach is necessary for their evaluation. This chapter provides a comprehensive overview of the modern computational framework used to tackle these challenges. We begin by establishing the fundamental definitions and the linear relations between integrals provided by Integration-By-Parts identities. These identities allow us to reduce a vast family of integrals to a minimal set of master integrals. The primary strategy for evaluating these master integrals involves the method of differential equations. By differentiating the integrals with respect to external kinematic invariants or internal masses, one can derive a system of coupled linear differential equations. Physical problems involving multiple mass scales frequently lead to non-trivial geometries where a solution in terms of known classes of functions is non trivial, or even not possible, to obtain. In such cases, numerical evaluation via series expansions becomes an indispensable tool. This approach allows for high-precision results across the entire phase space without having to deal with the analytic complexity of the integrals. The development of automated tools for series expansions is becoming extremely valuable for the community. A first implementation of this strategy was the Mathematica package DIFFEXP [106]. However, DIFFEXP was designed only to handle real mass parameters, meaning that it cannot be used to handle intermediate unstable particles such as W - and Z -bosons. Motivated by the need to overcome these constraints, this thesis introduces SEASYDE (Series Expansion Approach for SYstems of Differential Equations) [2]. SEASYDE represents an original contribution to the field, developed to extend the power of series expansion methods to the complex domain. By generalizing the algorithms for the analytic continuation, we are able to handle complex-valued parameters, offering a more robust framework for electroweak calculations. In Section 3.7 we present SEASYDE, while in Section 3.7.2 we present some benchmarks to validate it.

3.1 Definitions

The object that we are interested in computing are Feynman Integrals (FIs), i.e. dimensionally regularized integrals [83] of the following form:


$$I_{\alpha}(s_j; d) = \int \prod_{k=1}^l \frac{d^d q_k}{(2\pi)^d} \frac{1}{\mathcal{D}_1^{\alpha_1} \dots \mathcal{D}_n^{\alpha_n}}, \quad (3.1)$$

where l is the number of loops, $d = 4 - 2\epsilon$ is the number of dimensions, s_j are the kinematic invariants that the FIs depend on, q_k are the loop momenta and $\mathcal{D}_i = p_i^2 - m_i^2$ are inverse propagators. In particular, p_i is a linear combination of external and loop momenta while m_i is the mass of the particle running in the i -th propagator. Finally, each denominator is raised to an integer power α_i , which might also take negative values.

A given set of denominators $\{\mathcal{D}_i\}$ constitutes an integral family and within an integral family, each FI is uniquely identified by the set $\alpha = \{\alpha_1, \dots, \alpha_n\}$ of powers to which each denominator is raised. In the following, we will see that some scalar products might appear in the numerator of the integrand function. In order to bring again the integral in the form of Eq. 3.1, we have to be able to express each scalar product as a linear combination of denominators and kinematic invariants. Since we can have $l(l+1)/2 + l(n-1)$ independent scalar products, where n is the number of external legs, we need to have the same number of denominators. This means that we might have to define extra denominators, which are called auxiliary denominators. This is better understood with an example.

3.1.1 Example: two-loop massless triangle

We consider the following two-loop massless triangle:



with p_1, p_2 and p_3 the external momenta satisfying $p_1^2 = p_2^2 = p_3^2 = 0$ and $p_1 + p_2 = p_3$. By labeling q_1 and q_2 the upper loop propagators, we can fix all the others just by imposing conservation of momentum on each vertex. The list of six denominators then read:

$$\begin{aligned} \mathcal{D}_1 &= q_1^2, & \mathcal{D}_2 &= q_2^2, \\ \mathcal{D}_3 &= (q_1 - p_1)^2, & \mathcal{D}_4 &= (q_1 - q_2)^2, \\ \mathcal{D}_5 &= (q_1 - p_1 - p_2)^2, & \mathcal{D}_6 &= (q_2 - p_1 - p_2)^2. \end{aligned} \quad (3.2)$$

However, we have 7 possible scalar products, which are listed in Table 3.1. We introduce an auxiliary denominator $\mathcal{D}_7 = (q_2 - p_1)^2$, so that every scalar product can be written as a combination of $\mathcal{D}_i$. In particular, we find:

Scalar Product	Relation
q_1^2	$\mathcal{D}_1$
q_2^2	$\mathcal{D}_2$
$q_1 \cdot q_2$	$\frac{1}{2}(\mathcal{D}_1 + \mathcal{D}_2 - \mathcal{D}_4)$
$q_1 \cdot p_1$	$\frac{1}{2}(\mathcal{D}_1 - \mathcal{D}_3)$
$q_1 \cdot p_2$	$\frac{1}{2}(\mathcal{D}_3 - \mathcal{D}_5)$
$q_2 \cdot p_1$	$\frac{1}{2}(\mathcal{D}_2 - \mathcal{D}_7)$
$q_2 \cdot p_2$	$\frac{1}{2}(\mathcal{D}_7 - \mathcal{D}_6)$

Table 3.1: Relations between scalar products and denominators for the two-loop massless triangle

If, at any point of the calculation, we find a scalar product in the numerator, we can bring the FI back to the form of Eq. 3.1 using the relations in the table above. For example:

$$\begin{aligned}
\int \frac{d^d q_1}{(2\pi)^d} \frac{d^d q_2}{(2\pi)^d} \frac{q_1 \cdot q_2}{\mathcal{D}_1 \mathcal{D}_2 \mathcal{D}_4 \mathcal{D}_5} &= \frac{1}{2} \int \frac{d^d q_1}{(2\pi)^d} \frac{d^d q_2}{(2\pi)^d} \frac{\mathcal{D}_1 + \mathcal{D}_2 - \mathcal{D}_4}{\mathcal{D}_1 \mathcal{D}_2 \mathcal{D}_4 \mathcal{D}_5} = \\
&= \frac{1}{2} (I_{0,1,0,1,1,0,0} + I_{1,0,0,1,1,0,0} - I_{1,1,0,0,1,0,0}). \quad (3.3)
\end{aligned}$$

This will be useful when obtaining the differential equations.

3.2 Integration-by-parts identities

During state-of-the-art calculations, one usually has to compute up to hundred of thousands of FIs. However, not all of them are independent, but there exist linear relations among them. These relations go under the name of Integration-by-parts Identities (IBPs) and they follow from Gauss's theorem in d -dimensions:

$$\int \prod_{k=1}^l \frac{d^d q_k}{(2\pi)^d} \frac{\partial}{\partial q_i^\mu} \left\{ v_\mu \frac{1}{\mathcal{D}_1^{\alpha_1} \dots \mathcal{D}_n^{\alpha_n}} \right\} = 0, \quad (3.4)$$

where v_μ can be either a loop or an external momentum. The first step is to write down identities from Eq. 3.4, for different values of v_μ and for all the α that appear in our calculation. Then we can solve the system by Gauss substitution rule, i.e. considering each equation, expressing one integral in terms of the others and substitute it in the remaining ones. This procedure can be approached in a systematic way by using Laporta algorithm [107]. This is done by assigning a weight to each integral and by systematically expressing integrals of higher weight in terms of others with lower weight. Laporta algorithm is extremely simple and elegant and it has been implemented in several public codes, such as AIR [108], Blade [109], FIRE [110, 111], Kira [112–114], LiteRed [115], NeatIBP [116] and Reduze [117]. Its algebraic complexity, however, grows very fast with the number of loops, external legs and internal masses. For this reason, for state-of-the-art problems, Laporta algorithm is used in combination with a more advanced technique, the method of finite fields. The idea behind this method is to solve the system of IBPs numerically multiple times over finite fields and then, at the end, reconstruct the symbolic rational coefficients for the identities of interest by combining the samples together. The advantage of this ap-

proach is that the implementation of modular arithmetic in statically typed languages such as C or C++ is fast, exact and suitable for parallelization since each sample is independent from the others. The idea was originally introduced in [118], and later implemented in public packages such as FiniteFlow [119], FireFly [120] and Ratracer [121]. In general, obtaining compact IBPs in a reasonable amount of time is one of the bottleneck of the calculations.

After the implementation of IBPs we are able to express every integral, belonging to the a given integral family, in terms of a finite number of FIs which are called Master Integrals (MIs). These Master Integrals play the role of a basis in the space of Feynman integrals. Note that the term basis may be misleading, because a priori we do not know if this set of MIs is minimal, that is if there exist other relations between them and so if they are truly independent on from the other, anyway this term gives a good idea of what is going on.

3.3 Differential equations

In this section we will present a general method to evaluate the MIs introduced in the previous section: the method of differential equations, which was proposed in a simplified version by A. Kotikov and later improved by E. Remiddi and T. Gehrmann, who successfully applied this method to some significant examples [122–125].

3.3.1 The method

This method aims at writing down a system of differential equations with respect to one of the kinematic invariants of the problem, and, subsequently, solve it. The first step is to differentiate all the MIs w.r.t. one of the kinematic invariants. This is done by exploiting the chain rule:

$$\frac{\partial}{\partial s_k} = \sum_{\substack{i,j=1 \\ i \leq j}}^{n-1} \frac{\partial(p_i \cdot p_j)}{\partial s_k} \frac{\partial}{\partial(p_i \cdot p_j)}, \quad (3.5)$$

where n is the number of external momenta. The derivative of a MI w.r.t. the scalar product $(p_i \cdot p_j)$ can be computed with the following formulas:

$$\begin{aligned} \frac{\partial I_\alpha(\mathbf{s}; d)}{\partial(p_i \cdot p_j)} &= \sum_k [\mathbb{G}^{-1}]_{kj} p_k \cdot \frac{\partial I_\alpha(\mathbf{s}; d)}{\partial p_i} = \sum_k [\mathbb{G}^{-1}]_{ki} p_k \cdot \frac{\partial I_\alpha(\mathbf{s}; d)}{\partial p_j}, \\ \frac{\partial I_\alpha(\mathbf{s}; d)}{\partial p_i^2} &= \frac{1}{2} \sum_k [\mathbb{G}^{-1}]_{ki} p_k \cdot \frac{\partial I_\alpha(\mathbf{s}; d)}{\partial p_i}, \end{aligned} \quad (3.6)$$

where $\mathbb{G} = \mathbb{G}(p_1, \dots, p_{n-1})$ is the Gram matrix, i.e.

$$\mathbb{G}(p_1, \dots, p_{n-1}) = \begin{pmatrix} p_1^2 & p_1 \cdot p_2 & \dots & p_1 \cdot p_{n-1} \\ p_1 \cdot p_2 & p_2^2 & \dots & p_2 \cdot p_{n-1} \\ \vdots & & \ddots & \vdots \\ p_1 \cdot p_{n-1} & p_2 \cdot p_{n-1} & \dots & p_{n-1}^2 \end{pmatrix} \quad (3.7)$$

Note that when applying the operator in Eq. 3.5 to an integral of the form of Eq. 3.1, only two things can happen. The power α_j to which one inverse denominator is raised can increase by 1 and/or a scalar product of loop and external momenta can appear in the numerator. The latter is then reduced to a combination of $\mathcal{D}_i$ as shown in Sec. 3.1.1. During this procedure only denominators belonging to the

given integral family can appear, meaning that is possible to express the derivative of a MI in terms of a linear combination of other FIs belonging to the same integral family. Finally, we can use again IBPs identities to reduce this combination of FIs to a linear combination of MIs. By repeating this process for all the MIs we end up with a system of homogeneous first order linear differential equations, to which the MIs are a solution.

3.3.2 Example: 1L QED vertex

Let us consider the following integral family, which appears in the calculation of the 1-loop vertex correction in Quantum Electrodynamics (QED):

$$I_{\alpha_1, \alpha_2, \alpha_3}(s; d) = \int \frac{d^d q}{(2\pi)^d} \frac{1}{[q^2]^{\alpha_1} [(q+p_1)^2 - m^2]^{\alpha_2} [(q-p_2)^2 - m^2]^{\alpha_3}}. \quad (3.8)$$

This integral family has only two MIs, namely $I_{0,1,0}$ and $I_{1,1,1}$, two independent momenta, p_1 and p_2 , and one kinematic invariant $s = (p_1 + p_2)^2$. The Gram matrix reads:

$$\mathbb{G} = \begin{pmatrix} m^2 & \frac{s-2m^2}{2} \\ \frac{s-2m^2}{2} & m^2 \end{pmatrix}, \quad \mathbb{G}^{-1} = \frac{1}{s(4m^2 - s)} \begin{pmatrix} 4m^2 & 4m^2 - 2s \\ 4m^2 - 2s & 4m^2 \end{pmatrix}. \quad (3.9)$$

The first integral does not depend on s , so its differential equation is trivial. For the second one we have:

$$\frac{\partial I_{1,1,1}}{\partial s} = \frac{1}{2} \left[\frac{(p_1 + p_2)^\mu}{s} + \frac{(p_1 - p_2)^\mu}{s - 4m^2} \right] \frac{\partial I_{1,1,1}}{\partial p_1^\mu}. \quad (3.10)$$

The derivative w.r.t. p_1^μ in Eq. 3.10 reads:

$$\frac{\partial I_{1,1,1}}{\partial p_1^\mu} = \int \frac{d^d q}{(2\pi)^d} \frac{\partial}{\partial p_1^\mu} \left(\frac{1}{\mathcal{D}_1 \mathcal{D}_2 \mathcal{D}_3} \right) = \int \frac{d^d q}{(2\pi)^d} \frac{-2(q+p_1)^\mu}{\mathcal{D}_1 \mathcal{D}_2^2 \mathcal{D}_3}. \quad (3.11)$$

Now we can insert it in Eq. 3.10, use $q \cdot p_1 = (\mathcal{D}_2 - \mathcal{D}_1)/2$ and $q \cdot p_2 = (\mathcal{D}_1 - \mathcal{D}_3)/2$, and simplify the expression. In the end we obtain:

$$\frac{\partial I_{1,1,1}}{\partial s} = \frac{1}{s - 4m^2} I_{0,2,1} - \frac{s - 2m^2}{s(s - 4m^2)} I_{1,1,1} - \frac{2m^2}{s(s - 4m^2)} I_{1,2,0}. \quad (3.12)$$

Finally, we use again the IBPs in order to reduce the integral in the r.h.s. of Eq. 3.12 to a combination of the two MIs of the family:

$$\frac{\partial}{\partial s} \begin{pmatrix} I_{0,1,0} \\ I_{1,1,1} \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ \frac{\epsilon-1}{sm^2(s-4m^2)} & \frac{2m^2-s-\epsilon}{s(s-4m^2)} \end{pmatrix} \begin{pmatrix} I_{0,1,0} \\ I_{1,1,1} \end{pmatrix}. \quad (3.13)$$

Before moving on, a few comments are in order. The first one regards the singular structure of the MIs. In particular, we observe that this can be read from the poles of the coefficient matrix in Eq. 3.13. Indeed, we see that there are two of them, one for $s = 0$ and the other for $s = 4m^2$. The first one is a pseudo-threshold, while the second is the physical threshold associated to the on-shell production of the two internal fermions.

A second comment regards the number of variables that appear in the system. A common procedure is to introduce adimensional variables, defined as ratios of kinematical invariants. This is done in order to reduce by one the total number of variables in the differential equations, thus simplifying the solution procedure. For example, one could introduce the variable $x = s/m^2$, so that the system in Eq. 3.13

becomes

$$\frac{\partial}{\partial x} \begin{pmatrix} I_{1,0,0} \\ I_{1,1,1} \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ \frac{\epsilon-1}{x(x-4)} & \frac{2-x-x\epsilon}{x(x-4)} \end{pmatrix} \begin{pmatrix} I_{1,0,0} \\ I_{1,1,1} \end{pmatrix}. \quad (3.14)$$

3.3.3 Boundary conditions

Before diving into how to solve the system of differential equations let us address the problem of finding the boundary conditions. The boundary conditions, in general, could be given in three different forms. The first one is as the value of MIs in a particular point in the phase-space. In this case we can construct a well-defined Cauchy problem. In order to obtain the boundary conditions in this form we need to solve the MIs explicitly in that particular point of the phase-space, either analytically or numerically. If we choose to tackle the problem analytically, one usually chooses a point in the phase space in which most of the kinematic invariants vanish, so that the integrals are easier to evaluate. However, a Feynman integral usually develops divergences as we approach such a point, and hence by simply plugging these values into the integrand, we would not obtain the correct asymptotic limit. One possible solution to this problem is given by the method of expansion by regions [126–128]. Another possibility is to obtain them numerically in the Euclidean region. This can be done, for example, with PYSECDEC [129], which parametrises the integral and then perform the parametric integration with Monte-Carlo techniques. The problem of this approach is that the numerical precision is limited. Finally, one can use the Auxiliary Mass Flow method, implemented in the Mathematica package AMFLOW [130].

As alternative to boundary conditions, there are two other options for fixing the homogenous solution of the differential equations. It may happen, indeed, that from independent arguments we know that the integral must be regular at a particular point of the phase-space, usually a pseudo-threshold. In this cases, it might be possible to impose the finiteness of the solution to completely fix it. A final possibility is to fix the asymptotic behavior of the solution on a threshold. In particular, in some cases fixing the coefficient of the logarithmic term of the solution might be sufficient to completely determine it. The choice of a method or the other depends on the particular case we are considering.

3.3.4 Epsilon expansion

From now on, we will consider only one kinematic variable at the time. This choice will simplify the calculations and will be made clearer when discussing the analytic continuation of the solution. To this end, we suppose that the boundary conditions are imposed in $\tilde{\mathbf{s}} = \{\tilde{s}_1, \dots, \tilde{s}_m\}$ and that we are solving the system with respect to the variable s_j . In the system we perform the substitution

$$s_i \rightarrow \tilde{s}_i \quad \text{with } i \neq j \quad (3.15)$$

so we obtain a system with only one kinematic variable s_j . From this moment on, for simplicity, we will drop the index j . So the system, in general, takes the following form:

$$\frac{\partial}{\partial s} \vec{I}(s; \epsilon) = \mathbf{A}(s; \epsilon) \vec{I}(s; \epsilon), \quad (3.16)$$

where $\vec{I}(s; \epsilon)$ denotes the vector of MIs. First of all, the system needs to be ϵ -expanded. To this end let us write:

$$\vec{I}(s; \epsilon) = \sum_{k=\epsilon_{min}}^{\infty} \vec{I}^{(k)}(s) \epsilon^k, \quad (3.17)$$

here we assumed that the expansion starts at order ϵ_{min} . The value of ϵ_{min} can be determined from the boundary conditions. Along with $\vec{I}$ we expand also the coefficient matrix $\mathbf{A}$.

$$\mathbf{A}(s; \epsilon) = \sum_{k=0}^{\infty} \mathbf{A}^{(k)}(s) \epsilon^k \quad (3.18)$$

We can always assume that there are no poles in ϵ , since they can be removed by rescaling the basis $\vec{I}$ by an overall power of ϵ . Using this expansion in the Eq. 3.16 and collecting order-by-order the different powers in ϵ we obtain a tower of equations

$$\partial_s \vec{I}^{(k)}(s) = \mathbf{A}^{(0)}(s) \vec{I}^{(k)}(s) + \sum_{j=\epsilon_{min}}^{k-1} \mathbf{A}^{(k-j)}(s) \vec{I}^{(j)}(s). \quad (3.19)$$

An important thing to notice is that the equations for $\vec{I}^{(k)}$ involve $\vec{I}^{(j)}$ with $j < k$ but not with $j > k$. So we can start from $k = \epsilon_{min}$ and solve the system to obtain $\vec{I}^{(\epsilon_{min})}$. Now we substitute $\vec{I}^{(\epsilon_{min})}$ in the equation for $\vec{I}^{(\epsilon_{min}+1)}$ and solve it with respect to $\vec{I}^{(\epsilon_{min}+1)}$ and we proceed recursively up to the desired order in ϵ .

3.3.5 Canonical basis

Here we will discuss a completely general method to solve the system of differential equations. The idea is to perform a change of basis and to write down the solution of the system as a series expansion in ϵ . This method was first proposed by J. Henn [131], and here we will illustrate it briefly.

In the previous section we managed to write down a system of differential equations in the form of Eq. 3.16. We consider a change of basis $\vec{I} \rightarrow \mathbf{B} \vec{I}$. Under such a transformation it can be shown that the matrices $\mathbf{A}$ transforms

$$\mathbf{A} \rightarrow \mathbf{B}^{-1} \mathbf{A} \mathbf{B} - \mathbf{B}^{-1} \left(\frac{\partial}{\partial s} \mathbf{B} \right). \quad (3.20)$$

If $\mathbf{B}$ is chosen such that $\mathbf{A} \rightarrow \epsilon \tilde{\mathbf{A}}$, that is, if we could factorize an overall ϵ factor, the system is in the so-called canonical form. In [131] it has been conjectured that such a basis always exists, even though this has not been proved yet. The equations for a system in canonical form can be formally solved

$$\vec{I} = \mathbb{P} \exp \left(\epsilon \int_{\gamma} d\tilde{\mathbf{A}} \right) \vec{I}(\gamma(0)), \quad (3.21)$$

here $\gamma : [0, 1] \rightarrow \mathbb{C}^{|s|}$ is a path in the phase-space of the kinematic invariants $\mathbf{s}$, and $|s|$ denotes the number of these, while $\mathbb{P}$ denotes a path-ordered integral. Given this formal expression for $\vec{I}$ then the solution can be formally written as

$$\begin{aligned} \vec{I}(\mathbf{s}) = & \vec{I}^{(0)}(\gamma(0)) + \\ & + \sum_{k \geq 1} \epsilon^k \sum_{j=1}^k \int_0^1 \gamma^*(d)(t_1) \int_0^{t_1} \gamma^*(d)(t_2) \cdots \int_0^{t_{j-1}} \gamma^*(d)(t_j) \vec{I}^{(k-j)}(\gamma(0)) \end{aligned} \quad (3.22)$$

where γ^* is the pull-back of the differential forms to the unit interval $[0, 1]$. These iterated integrals are called *Chen iterated integrals*.

3.4 Series expansion approach

In this section we will discuss how to solve the system for each order in ϵ using the series expansion approach. The method was first introduced in Ref. [132] and was shortly after implemented in the public Mathematica package DIFFEXP [106]. The idea of this approach consists in expanding the equations both in ϵ and in a kinematic variable s , so that the problem of solving a differential equation reduces to an algebraic one. Its main advantage with respect to analytic ones, is that the complexity of the latter grows fast when increasing the number of external legs and/or the number of internal scales. While for the series expansion approach, since at each step of the calculation we are dealing with power series and logarithms, all the steps can be carried out analytically and in a systematic way, independently from the mathematical structure of the problem. Moreover, the numerical precision of the solution can be directly controlled by the number of terms that we consider, meaning that, provided that we have infinite time and space, we could achieve, in principle, arbitrary precision. Finally, once we have the solution, it can be evaluated numerically in a negligible amount of time. Power series, however, have some drawbacks. The biggest one being that series have a limited radius of convergence, meaning that we have to provide an algorithm for performing the analytic continuation of the result. This will be the focus of Sec. 3.5.

3.4.1 Bottom-to-top approach and solution of equations

Let us look closer at the system in Eq. 3.19, i.e. let us consider a fix order $\epsilon^{\bar{k}}$ and let us assume that we already solved all the ones for $k < \bar{k}$. This is a system of first order linear differential equations, in which the equations are in principle coupled. However, in many cases by choosing an appropriate base of MIs, the system in Eq. 3.19 can be casted in an upper-triangular form, i.e.

$$\frac{\partial}{\partial s} \begin{pmatrix} I_1^{(\bar{k})} \\ I_2^{(\bar{k})} \\ I_3^{(\bar{k})} \\ \vdots \\ I_n^{(\bar{k})} \end{pmatrix} = \begin{pmatrix} \star & \star & \star & \dots & \star \\ 0 & \star & \star & \dots & \star \\ 0 & 0 & \star & \dots & \star \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \star \end{pmatrix} \begin{pmatrix} I_1^{(\bar{k})} \\ I_2^{(\bar{k})} \\ I_3^{(\bar{k})} \\ \vdots \\ I_n^{(\bar{k})} \end{pmatrix} + \begin{pmatrix} \star \\ \star \\ \star \\ \vdots \\ \star \end{pmatrix}. \quad (3.23)$$

If the system is in the form of Eq. 3.23, it can be solved using a bottom-to-top approach. Indeed, the last equation contains only $I_n^{(\bar{k})}$, so we can solve for it and substitute the solution into the second to last equation, solve for $I_{n-1}^{(\bar{k})}$ and proceed recursively until we solve the system completely.

At each step of the process we have to solve a first order linear differential equation, with a given boundary condition.

$$\begin{cases} \frac{\partial}{\partial s} f(s) = a(s)f(s) + h(s), \\ f(s_0) = \tilde{f} \end{cases} \quad (3.24)$$

where for simplicity we substituted $I_i^{(\bar{k})} \equiv f$ and $a(s)$ and $h(s)$ are rational functions of s . For solving Eq. 3.24 we can proceed by solving the homogeneous equation, finding a particular solution and, lastly, by imposing the boundary condition.

3.4.2 Frobenius method

In order to solve the homogeneous equation we use the Frobenius method. This is a general technique for solving a homogeneous ordinary differential equation

$$\frac{\partial}{\partial s} f(s) = a(s)f(s) \quad (3.25)$$

around a point $s = s_0$. Without loss of generality from now on we consider $s_0 = 0$. Let us consider the following ansatz for the solution:

$$f_{hom}(s) = s^r \sum_{m=0}^{\infty} c_m s^m \quad (3.26)$$

where c_m are complex coefficients and r is a rational exponent. The term s^r has the role to take into account for some divergent terms in the solution. We series expand the function $a(s)$ in the differential equations, then by substituting Eq. 3.26 into Eq. 3.25, and collecting the different terms based on different powers of s , we obtain a set of algebraic equations for the coefficients c_m . At leading order in s , the equation is a non-trivial polynomial equation for r , which is called indicial equation. After fixing r we may (recursively) solve for all c_m with $m \geq 1$. In the end there is one last free parameter, namely c_0 . An example of application of the Frobenius method is given in 3.4.4.

3.4.3 Variation of parameters

Once we have the solution for the homogeneous equation, we can obtain a particular solution using the variation of parameters method. We consider the following ansatz for a particular solution:

$$f_{part}(s) = C(s)f_{hom}(s) \quad (3.27)$$

where $C(s)$ is a function to be determined. Let us substitute Eq. 3.27 into Eq. 3.24

$$C'(s)f_{hom}(s) + C(s)f'_{hom}(s) = a(s)C(s)f_{hom}(s) + h(s) \quad (3.28)$$

but from Eq. 3.25 we see that $f'_{hom}(s) = a(s)f_{hom}(s)$ hence we are left with:

$$C'(s)f_{hom}(s) = h(s). \quad (3.29)$$

Finally we invert it to obtain $f_{part}(s)$:

$$f_{part}(s) = f_{hom}(s) \int_{\bar{s}}^s h(s') f_{hom}^{-1}(s') ds'. \quad (3.30)$$

Now that we have a particular solution we can construct the general one by summing the homogeneous together with the particular one:

$$f(s) = c f_{hom}(s) + f_{part}(s) \quad (3.31)$$

where c is a complex constant that is determined imposing the boundary condition. Note that since f_{hom} is a series, after expanding $h(s)$, the product $h f_{hom}^{-1}$ is still a series and hence it can be integrated analytically.

3.4.4 Example: 1L QED Vertex

Let us illustrate how this works in practice by considering the order $1/\epsilon$ of the system for the 1-loop QED Vertex, given in Eq. 3.14:

$$\begin{cases} \frac{\partial}{\partial x} B_1^{(-1)} = 0 \\ \frac{\partial}{\partial x} B_2^{(-1)} = -\frac{1}{x(x-4)} B_1^{(-1)} - \frac{x-2}{x(x-4)} B_2^{(-1)} \\ B_1^{(-1)}(0) = 1 \\ B_2^{(-1)}(0) = 1/2 \end{cases} \quad (3.32)$$

where for readability we defined $B_1 \equiv I_{0,1,0}$ and $B_2 \equiv I_{1,1,1}$. The system is indeed in a triangular form. In particular, the first equation is trivial and gives $B_1^{(-1)} = 1$. We start solving the second one by considering the ansatz

$$B_2^{(-1),hom}(x) = x^r \sum_{i=0}^{\infty} c_i x^i \quad (3.33)$$

Substituting in the homogeneous equation and collecting different powers of x we find:

$$\begin{cases} r c_0 = -\frac{1}{2} c_0 \\ (1+r) c_1 = \frac{1}{8} - \frac{c_1}{2} \\ (2+r) c_2 = \frac{1}{32} (1 + 4c_1 - 16c_2) \\ (3+r) c_3 = \frac{1}{128} (1 + 4c_1 + 16c_2 - 64c_3) \\ \dots \end{cases} \quad (3.34)$$

The indicial equation gives $r = -1/2$ and then by recursively solving for all the c_i we find:

$$B_2^{(-1),hom}(x) = \frac{c_0}{\sqrt{x}} \left(1 + \frac{1}{8}x - \frac{3}{128}x^2 + \frac{5}{1024}x^3 + \mathcal{O}(x^4) \right). \quad (3.35)$$

In order to find the particular solution we use Eq. 3.30:

$$B_2^{(-1),part}(x) = \frac{1}{2} + \frac{x}{12} + \frac{x^2}{60} + \frac{x^3}{280} + \mathcal{O}(x^4) \quad (3.36)$$

The complete solution is hence:

$$B_2^{(-1)}(x) = B_2^{(-1),hom}(x) + B_2^{(-1),part}(x) = \quad (3.37)$$

$$\begin{aligned} &= c_0 x^{-1/2} \left(1 + \frac{x}{8} + \frac{3x^2}{128} + \frac{5x^3}{1024} + \mathcal{O}(x^4) \right) + \\ &\quad + \left(\frac{1}{2} + \frac{x}{12} + \frac{x^2}{60} + \frac{x^3}{280} + \mathcal{O}(x^4) \right) \end{aligned} \quad (3.38)$$

Finally, we fix the boundary condition $c_0 = 0$.

Note that, in this case knowing that the solution was regular for $x = 0$ would have been sufficient for fixing the boundary condition. From the coefficient matrix we read that there are two possible singular points, i.e. $x = 0$ and $x = 4$. We just found out that the solution for $x = 0$ is regular, meaning that $x = 0$ is a pseudo-threshold. We could have also chosen to solve the system around a non problematic point, like for example $x = 1$. In this case, we are guaranteed that the solution is always a simple Taylor series. Lastly, we could have solved the equation on top of a physical threshold, like $x = 4$. In this case




Figure 3.1: Comparison of the series solution for $B_2^{(-1)}$ against the exact solution. The series is centered in $x = 0$ and the closest singularity is $x = 4$. This implies that the series converges in the interval $(-4, 4)$, which can be seen from the divergence of the solution around $x = -4$.

the final solution would have contained terms like

$$\frac{1}{x-4}, \quad \log(x-4) \quad (3.39)$$

which can arise from the integration in Eq. 3.30 or, at higher orders in ϵ , from the non-homogeneous term in the equation.

Finally a comment regarding the convergence of the series centered in $x = 0$. In particular, the radius of convergence is given by the distance between the center of the series with the closest singularity. In this case the radius of convergence is 4, which means that the series converges in the interval $(-4, 4)$. This can be seen explicitly in Fig. 3.1.

3.4.5 Coupled equations

In many practical applications, casting the system of differential equations in triangular form as in Eq. 3.23 requires extensive algebraic manipulations or the identification of a highly non-trivial basis of master integrals. Furthermore, if the underlying geometry of the problem features an elliptic sector, achieving such a fully decoupled triangular form becomes strictly impossible. In these scenarios, directly addressing the system in its coupled form is not just a more efficient alternative, but often the only viable option. In this section, we will discuss how to generalize the methods introduced previously to solve a system of coupled differential equations. Let us start by considering the following system of coupled equations

$$\left\{ \begin{array}{l} B_1'(x) = \frac{B_1(x)}{x} - \frac{3B_2(x)}{x} + \left(\frac{1}{2} - \frac{9}{x}\right) \\ B_2'(x) = -\frac{2(x-3)B_1(x)}{x(x-9)(x-1)} + \frac{2(5x-9)B_2(x)}{x(x-9)(x-1)} - \frac{648 + (4\pi^2 - 273)x + 27x^2}{12x(x-9)(x-1)} \\ B_1(1) = \frac{59}{8} + \frac{\pi^2}{4} \\ B_2(1) = \frac{\pi^2}{12} - \frac{1}{2} \end{array} \right., \quad (3.40)$$

and we try to solve it through an expansion around $x = 1$. In order to do so we consider the following ansatz for the solution of the homogeneous equation:

$$\begin{aligned} B_1^{hom}(x) &= (x-1)^r \sum_{i=0}^{\infty} a_i (x-1)^i \\ B_2^{hom}(x) &= (x-1)^r \sum_{i=0}^{\infty} b_i (x-1)^i. \end{aligned} \quad (3.41)$$

Now we proceed in the same way as in the single equation case, that is we substitute the ansatz in the homogeneous equation, we expand everything around $x = 1$, collect the different powers of $(x-1)$ and, finally, solve recursively the infinite set of algebraic equations for a_i and b_i . The result is:

$$\begin{aligned} B_1^{hom}(x) &= a_1 \left((x-1)^2 - \frac{5(x-1)^3}{4} + \frac{87(x-1)^4}{64} + \mathcal{O}(x-1)^5 \right) \\ B_2^{hom}(x) &= a_1 \left(-\frac{2(x-1)}{3} + \frac{11(x-1)^2}{12} - \frac{47(x-1)^3}{48} + \right. \\ &\quad \left. + \frac{97(x-1)^4}{96} + \mathcal{O}(x-1)^5 \right). \end{aligned} \quad (3.42)$$

By looking closely to the solution in Eq. 3.42 we see that it depends only on one parameter, namely a_1 . However, this is a system of two differential equations and, as such, we expect two linearly independent solutions. This is related to the fact that the ansatz we chose was not general enough. We replace, hence, the ansatz in Eq. 3.41 with a more general one:

$$\begin{aligned} B_1^{hom}(x) &= (x-1)^r \sum_{i=0}^{\infty} a_i (x-1)^i + \log(x-1) (x-1)^r \sum_{i=0}^{\infty} c_i (x-1)^i \\ B_2^{hom}(x) &= (x-1)^r \sum_{i=0}^{\infty} b_i (x-1)^i + \log(x-1) (x-1)^r \sum_{i=0}^{\infty} d_i (x-1)^i \end{aligned} \quad (3.43)$$

which contains also logarithmic terms. By proceeding always in the same way we find:

$$\begin{aligned} B_1^{hom}(x) &= a_0 \left[1 - \frac{x-1}{2} + \frac{9(x-1)^3}{128} + \mathcal{O}(x-1)^4 + \right. \\ &\quad \left. + \left(\frac{3(x-1)^2}{16} - \frac{15(x-1)^3}{64} + \mathcal{O}(x-1)^4 \right) \log(x-1) \right] + \\ &\quad + a_2 \left((x-1)^2 - \frac{5(x-1)^3}{4} + \mathcal{O}(x-1)^4 \right); \\ B_2^{hom}(x) &= a_0 \left[\frac{1}{2} - \frac{x-1}{16} - \frac{7(x-1)^2}{128} + \frac{71(x-1)^3}{1024} + \mathcal{O}(x-1)^4 + \right. \\ &\quad \left. + \left(-\frac{x-1}{8} + \frac{11(x-1)^2}{64} - \frac{47(x-1)^3}{256} + \mathcal{O}(x-1)^4 \right) \log(x-1) \right] + \\ &\quad + a_2 \left(-\frac{2(x-1)}{3} + \frac{11(x-1)^2}{12} - \frac{47(x-1)^3}{48} + \mathcal{O}(x-1)^4 \right) \end{aligned} \quad (3.44)$$

Now we see that the solution correctly depends on two different parameters, namely a_0 and a_2 . Thus for a system of n homogeneous differential equations we can use the general ansatz:

$$\vec{B}^{hom}(x) = (x - x_0)^r \sum_{i=0}^{\infty} \vec{c}_{i,0} (x - x_0)^i + (x - x_0)^r \sum_{j=1}^m \log^j(x - x_0) \sum_{i=0}^{\infty} \vec{c}_{i,j} (x - x_0)^i. \quad (3.45)$$

Firstly, we try to solve the system for $m = 0$. If we obtain n linearly independent solutions we are done, otherwise we increase the value of m , eventually up to $n - 1$. In principle, the ansatz in Eq. 3.45 is not strictly limited to polynomials. Indeed, any complete basis of functions could be chosen, as long as it correctly parametrizes the singular behavior of the solution in the vicinity of the expansion point x_0 . For example, when expanding around a regular point where the solution is entirely smooth, a basis of sines and cosines would be a mathematically valid alternative. However, because we are specifically targeting the local behavior around regular singular points, the generalized power-series ansatz is practically required to capture the fractional powers and logarithms while yielding tractable algebraic recurrence relations for the coefficients.

3.4.6 Variation of parameters for systems

The generalization of the variation of parameters technique to systems of differential equations is straightforward. Firstly, we organize the solution of the homogeneous equation in a matrix form:

$$\vec{B}^{hom}(x) = \mathbf{M}(x) \vec{a} \quad (3.46)$$

where $\mathbf{M}_{ij}(x)$ is the solution for the i -th MI where we put all the free parameters to 0, except for the j -th. And $\vec{a}$ is the vector of free parameters. Then a particular solution is given by:

$$\vec{B}^{part}(x) = \mathbf{M}(x) \int_{x_0}^x \mathbf{M}^{-1}(x') \vec{g}(x') dx' \quad (3.47)$$

where $\vec{g}(x')$ is the vector of non homogeneous terms of the equation. By following this procedure for the example introduced in the previous section we obtain:

$$\begin{aligned} B_1(x) &= \frac{59}{8} + \frac{\pi^2}{4} + \frac{3}{8}(x-1) + \frac{1}{4}(x-1)^2 - \frac{1}{3}(x-1)^3 + \mathcal{O}(x-1)^4 \\ B_2(x) &= \frac{1}{2} \left(\frac{\pi^2}{6} - 1 \right) + \frac{1}{4}(x-1)^2 - \frac{25}{96}(x-1)^3 + \mathcal{O}(x-1)^4. \end{aligned} \quad (3.48)$$

A few comments are in order. The first one is that within the integration in Eq. 3.47 only terms like $(x - x_0)^p \log^q(x - x_0)$ can appear. In practice, this integration can be performed with the implementation of substitution rules in order to make it faster. A second comment regards the inversion of the matrix $\mathbf{M}^{-1}(x)$. Inverting a matrix is an operation whose computational cost grows cubically with its dimension. For this reason, even if it might not be possible to write down a system in triangular form, it might be worth to cast it in a block triangular form.

3.5 Analytic continuation

Series have a limited radius of convergence, which is given by the distance of its center to the closest singularity. In order to extend the solution outside the region of convergence, we need to provide an algorithm for performing the analytic continuation in the entire complex kinematical plane. This can




Figure 3.2: Radius of convergence and example analytic continuation in the complex plane of the variable z .




Figure 3.3: Effect of a branch cut on the analytic continuation in the complex plane of the variable z . It shrinks the area in which the series converges to the desired value (left) and it modifies the path for moving from one point to another (right).

be seen graphically in Fig. 3.2. In particular, we are considering a complex variable z and we are discussing the analytic continuation in the z complex plane. The series is centered in z_0 and the radius of convergence is $\rho = \min_{w \in \mathcal{W}} |z - w|$, with $\mathcal{W} = \{w_0, w_+, w_-\}$ the set of all singularities. In the example shown in Fig. 3.2 the closest singularity is w_0 , hence, the series converges in the Γ_0 region. For extending the solution outside the blue region we have to perform an analytic continuation. This is done by evaluating the solution in a point within Γ_0 , for example z_1 , and then by solving again the system, this time centered around z_1 , and using the value we have just computed as a new boundary condition. It is possible to show that this procedure is unique. We saw also in the previous sections, that the solution might exhibit a logarithmic behavior. This complicates the analytic continuation procedure since it makes the solution a multi-valued function. In order to make it single-valued we have to introduce some branch-cuts, thus specifying which is the physical Riemann sheet on which we evaluate the solution. We decided to associate to each logarithmic singularity a branch-cut that starts from the singularity and goes to $-\infty$, parallel to the real axis. While this branch-cut does not modify the region in which the series converges, it shrinks the area in which the solution converges to the desired value, that is the one on the physical Riemann sheet. Looking at left panel of Fig. 3.3, while the series still converges in Γ_0 , only in the blue




Figure 3.4: Example of a Logarithmic expansion on the left, against a Taylor expansion on the right.

strip it does to the desired value. As a consequence, for extending the solution we have to choose a path that does not cross the branch cut. This is shown in the right panel of Fig. 3.3. For reaching z_1 we cannot go straight along the dashed line, because we would cross the branch cut and end up on the non physical Riemann sheet. Instead, we should follow the path of the solid red line, that avoids w_+ on the right. For readability of the picture, we have not plotted all the intermediate steps.

For following the red line, actually, there are two possibilities. In particular, they differ in the way they treat the singular point, either avoiding it or expanding on top of it. The two strategies are shown in Fig. 3.4 and, while they lead to the same result, they each have pros and cons. If we avoid the singularities (right panel of Fig. 3.4), at each step, we deal only with Taylor series, hence the solution is simpler and usually quicker to obtain. We dub this strategy Taylor expansion. The second possibility is the logarithmic expansion (left panel of Fig. 3.4). In this case, the solution might be slower to obtain, however, if the series is centered in w_+ the closest singularity is w_- , so the radius of convergence is bigger and we can take less intermediate steps. Choosing one strategies over the other is done case by case mainly by looking at the position of the other singularities which are present in the problem.

Finally, we discuss how to move along a branch cut. There are two different cases which are shown in Fig. 3.5. If there is not a singularity between the starting and ending point we can move straight since, by definition, we are never crossing a branch cut. If, on the contrary, we have a singularity between the two points we have to avoid moving either in the upper or lower part of the complex plane. This ambiguity is solved by considering Feynman prescription associated to the kinematical variable. A prescription $+i\delta$ implies an horse-shoe path in the upper half of the complex plane, while a $-i\delta$ in the lower one.

3.6 Mass expansion

The series expansion approach can be time consuming in state-of-the-art calculations. The running time, in particular, depends on different factors such as the number of MIs, the number of kinematic invariants which in turn increases the algebraic complexity of the differential equations, the point in which the BCs are imposed, the required precision, ... For this reason, a direct implementation in Monte-Carlo generators is, currently, out of reach. A common solution to the problem is to rely on pre-computed numerical grids for the complete amplitude. Their evaluation requires a big amount of time upstream, but then the interpolation in any point of the phase space is instantaneous and, hence, suitable for a




Figure 3.5: Examples of two possible ways to move along a branch cut. In case there is no a singularity between them (left) or there is (right).


Monte-Carlo implementation. Moreover, since the final corrections is a smooth function of the kinematic invariants, the numerical error on the interpolation is negligible.

The adoption of grids, however, comes with a severe limiting factor: all the input parameters are strictly fixed to their numerical values and cannot be modified without rerunning the final steps of the calculation. Nevertheless, the series expansion approach can be exploited to provide, with limited additional computational work, the continuous dependence on an extra variable, such as an internal mass m . This is achieved, firstly, by formulating the differential equations with respect to m and, secondly, by solving them using the original grid evaluated at a reference mass $\bar{m}$ as the boundary condition. In practice, this expansion procedure is performed directly at the level of the MIs, yielding a grid where each MI is represented as a power series in $\delta m = m - \bar{m}$. Since the full amplitude is simply a linear combination of these MIs, the corresponding rational coefficients, which depend analytically on the kinematics and the mass, can be straightforwardly Taylor-expanded in δm through standard algebraic operations. Multiplying the expanded coefficients by the expanded MIs systematically provides the complete amplitude as a local series in δm . This approach is particularly powerful, for example, in the W -boson mass determination where one has to generate many distributions for different values of m_W and then the best value for the W -boson mass is extracted with a template fitting procedure. Moreover, since the δm we are interested is usually small, only a small number of terms in the series is needed. Therefore, the solution can be computed in a small amount of time.

In order to illustrate this idea, we consider the following 2-loop box integral:

$$\mathcal{I} = \int \frac{d^d q_1}{(2\pi)^d} \frac{d^d q_2}{(2\pi)^d} \frac{1}{q_1^2 (q_2^2 - \mu_W^2) (q_1 - q_2)^2 (q_1 - p_1)^2} \cdot \frac{1}{(q_1 - p_1 - p_2)^2 (q_2 - p_1 - p_2)^2 ((q_2 - p_3)^2 - m_\ell^2)}, \quad (3.49)$$

which can be represented graphically as






Figure 3.6: On the left panel we plot the real and imaginary part of the $\mathcal{O}(\varepsilon^0)$ of $\mathcal{I}$ for different values of δm_W . In the right panel, we show the relative error in the $\mathcal{O}(\varepsilon^0)$ coefficient of $\mathcal{I}$, as a function of the expansion parameter δm_W . The error is evaluated as the absolute difference from unity of the ratio between the truncated series and the exact solution, obtained with AMFLOW.

Here the blue (red) indicates that an internal W -boson with mass μ_W (a final state lepton with mass m_ℓ) is running in the line. Such a diagram appears, for example, in the 2-loop mixed corrections for the Charged Current Drell-Yan. In the left panel of Figure 3.6 we plot the real and imaginary finite part of the integral, for different values of m_W , as a function of $\sqrt{s}$, with $\cos \theta = 0.165$. The intensity of the color corresponds to different choices of δm_W , from -500 MeV to 500 MeV, in steps of 250 MeV. The shift of the peak position illustrates the sensitivity of this integral to the choice of the mass value. In the right panel of Figure 3.6 we compare the exact evaluation of the integral¹ against different approximations. We consider $\sqrt{s} = 80.1315$ GeV in order to have a strong dependence on m_W and $\cos \theta = 0.165$. The curves correspond to a different number of terms in the δm_W expansion, and their ratio with the exact result is plotted as a function of δm_W . From the plot we can see that if we consider only shifts in δm_W of order 100 MeV, 15 terms in the expansions are sufficient for maintaining a relative precision of 10^{-15} .

3.7 SeaSyde

The theoretical frameworks and algorithms detailed in the previous sections, i.e. the generalized series expansion approach and the analytic continuation of the result to the entire (complex)-phase-space, have been integrated into the original Mathematica package SEASYDE (Series Expansion Approach for SYstems of Differential Equations). The package is publicly available and can be accessed via its official GitHub repository:

<https://github.com/TommasoArmadillo/SeaSyde>

The repository provides the source code, the documentation, and a selection of example notebooks that demonstrate the evaluation of both toy models and physical state-of-the-art multi-loop integrals. SEASYDE is built with a high degree of automation and ease of use in mind. The series expansion approach relies on the same generalized Frobenius ansatz to expand the solutions around any point of the phase-space, either singular or regular, entirely eliminating the need for process-specific calculations

¹The exact solution is obtained using AMFLOW, asking for a precision of 30 digits.

or tailored functional ansätze. The user is only required to provide the system of differential equations and the corresponding boundary conditions. From there, the package operates completely automatically: it systematically locates all singularities by analyzing the denominators in the matrix coefficients, solves the system by applying the universal ansatz, and performs the analytic continuation to the desired point in the (complex)-phase-space. A key strength of the package lies in its automated pathfinding capabilities. In particular, by implementing the procedure discussed in Section 3.5, SEASYDE automatically determines a viable integration path through the complex phase space to connect the boundary conditions to the target kinematic point. Alongside this automation, the package provides high flexibility by allowing the user to freely adjust algorithmic parameters, such as the truncation order of the series expansion and the step length of the analytic continuation. This ensures complete control over the numerical precision and computational performance of the evaluation. As an optional feature, users can also manually define custom integration paths. While this is not required for standard evaluations, it provides a valuable tool for theoretical exploration, allowing the user to investigate for example the behavior of the integrals when intentionally crossing branch cuts.

Crucially, the series expansion approach implemented in SEASYDE is completely general. It is capable of automatically tackling any state-of-the-art problem in particle physics, including those involving multiple scales and elliptic integrals, without requiring any theoretical extensions or structural improvements to the algorithm. The applicability of the method is inherently universal, the only limitation is purely computational. As the size of the differential equation system grows, the evaluation naturally becomes more demanding in terms of memory and execution time, but the underlying automated framework remains perfectly valid and intact.

3.7.1 Comparison with DIFFEXP

Here we compare SEASYDE with DIFFEXP [106], the first implementation of the series expansion approach, which was developed by Martijn Hidding in 2020. Both packages aim to solve systems of differential equations for master integrals by expanding the solutions in series and analytically continuing them across the phase space. However, a fundamental architectural difference distinguishes the two codes: their respective approaches to the analytic continuation algorithm and their treatment of kinematic singularities. DiffExp is designed to operate strictly with real-valued variables. Consequently, its analytic continuation algorithm is entirely confined to the real axis. Because the one-dimensional integration path, DIFFEXP cannot detour around these points and it is forced to expand the system exactly on top of the singular point using the full generalized Frobenius ansatz and then resume the continuation on the other side. This requirement means that the computationally heavy machinery of singular series expansions must be employed at every intermediate singularity encountered along the evaluation path. In contrast, SEASYDE treats all kinematic variables as intrinsically complex-valued quantities. By extending the kinematics into the complex plane, the package naturally gains an additional spatial dimension to work with. This extra degree of freedom fundamentally alters the analytic continuation strategy and allows SEASYDE to circumvent intermediate singularities entirely. Instead of colliding with a singular point on the real axis, the pathfinding algorithm simply deforms the contour, moving the integration path "up" or "down" into the complex plane to safely avoid the pole. Because the deformed path stays strictly within the regular regions of the differential equations, the analytic continuation from step to step can rely exclusively on standard, highly efficient Taylor series expansions. This architectural distinction strictly dictates the physical applicability of each package. Because DiffExp is restricted to real variables, it is well-suited for higher-order QCD calculations, where internal masses are purely real. However, it cannot be applied to the EW sector, which intrinsically features unstable particles that must be rigorously

described using the Complex Mass Scheme. Since SEASYDE natively supports complex-valued kinematics and accommodates arbitrary complex masses by design, it is the only viable tool between the two capable of performing exact multi-loop EW calculations.

3.7.2 Checks and performance

Here we present the most significant checks we performed to validate our implementation. Namely, the MIs which appear in the calculation of the two-loop mixed QCD-EW corrections of Neutral-Current Drell-Yan. The definitions and the details of the calculation will be given in Sec. 5.1, here we only report some highlights concerning the numerical evaluation of the MIs.

The most complicated topology is a two-loop planar box with two equal complex massive lines. The basis of masters is composed by 36 integrals that have to be evaluated. Out of these 36, for 31 of them an analytical expression in terms of Generalized Polylogarithms (GPLs) is publicly available [133–140], meaning that we can perform cross-checks using either real or complex masses. For the last 5, an analytical expression in terms of GPLs is not available, but only using Chen-Goncharov integrals [141]. In our case, the evaluation of the Chen iterated integrals in a non-physical region, where the boundary conditions are fixed, is possible, but the analytic continuation to cover the full physical region constitutes a formidable task. The need for the evaluation of these 5 MIs using complex internal masses was the main motivation behind the development of SEASYDE.

Concerning the first 31 integrals, we cross-checked the results given by SEASYDE against the analytic expressions evaluated with GINAC and HANDYG [142] in two independent C++ and Mathematica codes, in different point of the phase-space. We also use HarmonicSums [143, 144], PolyLogTools [145] and LoopTools [146] for several cross-checks of the one- and two-loop integrals appearing in the amplitudes. All these checks were performed with both real and complex masses.

For the last 5 MIs we validated our numerical results against DIFFEXP (restricted to real masses) and AMFLOW (using both real and complex masses). To quantify the level of agreement in the real-mass limit, we evaluated these integrals at multiple kinematic points across the physical phase space. Using 150 terms in the series expansion we found perfect numerical coincidence between SEASYDE, DIFFEXP, and AMFLOW up to the 40th decimal digit, corresponding to quadruple floating-point precision. Furthermore, we compared our results against PYSECDEC [129], though this was limited to a few points in the Euclidean region and six significant digits due to the nature of the sector decomposition approach. Notably, while the real-mass evaluations serve as a robust cross-check, the evaluation of these final 5 MIs with internal complex masses represents a genuinely new prediction achieved by SEASYDE.

The core strength of the series expansion approach is the ability to achieve an arbitrary level of theoretical precision simply by increasing the number of terms in the expansion. However, this increase in precision must be carefully balanced against the corresponding increase in computational time and resource consumption. To systematically benchmark this, we evaluated the MIs at various phase-space points using different expansion depths. As previously mentioned, a 150-term expansion yields 40 digits of precision. In Figure 3.7 (left), we illustrate this validation for one of the 31 MIs. The continuous lines generated by SEASYDE perfectly overlap with the independent GINAC evaluation, denoted by crosses.

In the plot we can also appreciate the effect of having a complex internal mass instead of a real one. The masters with real-valued masses exhibit a non differentiable behavior around the internal threshold $\sqrt{s} = 2m_W$, while they become smooth when moving to the Complex Mass Scheme, making the latter particularly important for phenomenological studies in the threshold region. The definition of the complex mass as $\mu^2 = m^2 - i\Gamma m$, shows that the case of real masses can be thought also as the limit for $\Gamma \rightarrow 0$. Hence, by considering values of Γ gradually smaller, all the MIs should converge to the




Figure 3.7: In the left panel, we see the real (solid) and imaginary (dashed) part of one the 31 MIs, computed both with real (blue) and complex (orange) masses. The continuous lines have been obtained using SEASYDE while the crosses have been independently computed with GiNaC. This serves as a cross-check of our implementation. In the right panel we see the real (solid) and imaginary (dashed) part of master 35 computed for different values of the decay width Γ . We see that for $\Gamma \rightarrow 0$ we recover the real mass case, as expected.

one obtained with real-valued internal masses. This is shown explicitly for one of the MIs in Figure 3.7 right, although it holds for all the other MIs. In particular, we have compared the results for real- and complex-valued masses, where the latter have been calculated using different values of the decay-width Γ . This example shows that SEASYDE can robustly handle arbitrary internal complex masses.

Regarding the time performances, they strongly depend not only on the number of terms in the series, but also on the presence of singularities between the starting and ending points. In Table 3.2 we report the time necessary for transporting the boundary conditions from the Euclidean region to a test point for different number of terms, along with the precision of the result.

Number of terms	Precision	Execution time
50 terms	10^{-14}	~ 14 min
75 terms	10^{-19}	~ 26 min
100 terms	10^{-25}	~ 50 min
125 terms	10^{-33}	~ 75 min
150 terms	10^{-40}	~ 90 min

Table 3.2: Precision achieved and execution time for transporting the boundary conditions from the Euclidean region to a test point for different number of terms in the series for the two-loop planar box topology with two massive lines.

This estimate is obtained on a MacBook Pro (2015) equipped with a 2.7 GHz Dual-Core Intel i5. The fact that the execution time increases so much with the number of terms, is mainly due to the poor performances of the Mathematica function `Series` when dealing with series with a big number of terms. Anyhow, for most current precision studies, a 10^{-19} precision is sufficient. The execution time depends also on the distance between the starting point, where the boundary conditions are imposed, and the final one, in which we would like to evaluate the solution. For reaching the physical region we usually need a number between 15 and 20 steps. Given the results in Table 3.2 we can estimate a total of 50, 160 and 300 s per step for 50, 100 and 150 terms, respectively. Once in the physical region, less singularities

need to be avoided and, hence, more direct path can be chosen. In this case a number between 1 and 10 steps is sufficient to reach every point in the phase-space. Having a BC directly in the physical region enormously speeds up the computation.

From amplitudes to physical observables

“On this proud and beautiful mountain we have lived hours of fraternal, loving, and exhausting struggle. Here, we cease to be the slaves of our own imagination.”

Lionel Terray, *French alpinist*.

In Chapter 2 and 3 we explored the computation of scattering amplitudes and the evaluation of Feynman integrals within the framework of perturbative QFT. These calculations are formulated in terms of fundamental, non-interacting initial and final states, namely quarks, gluons, and leptons. However, at modern colliders such as the LHC, the colliding beams consist of protons. A proton is a complex, bound state of partons (quarks and gluons) characterized by a low momentum scale, $\Lambda_{\text{QCD}} \approx 200 \text{ MeV}$, where the strong coupling constant α_s becomes $\mathcal{O}(1)$. Consequently, the interactions binding the proton together fall into the non-perturbative regime of QCD, rendering a purely perturbative calculation of a proton-proton scattering cross-section impossible. To make theoretically rigorous predictions for hadronic observables, we must isolate the high-energy, short-distance perturbative physics (which we can calculate using the techniques from previous chapters) from the low-energy, long-distance non-perturbative physics. This separation is rigorously formalized by the factorization theorem. In this Chapter we will first review the factorization theorem, we will sketch out how the calculation of Next-to-Leading Order (NLO) and Next-to-Next-to-Leading Order (NNLO) corrections works, focusing specifically on the Drell-Yan processes, and, finally we will introduce OCEANN [6] which is the framework we developed for performing the calculation of two-loop amplitudes.

4.1 Factorization theorem

The essence of the factorization theorem is that the timescale of the hard, high-momentum-transfer collision, $1/Q$, is vastly shorter than the timescale of the interactions binding the partons within the hadron, $1/\Lambda_{\text{QCD}}$. Thus, the incident partons can be treated as essentially free during the hard scattering.

For an inclusive hadronic collision process $A+B \rightarrow f+X$, where A and B are the incoming hadrons, f is a generic final state with a characteristic hard scale Q and X accounts for any additional radiation. By remaining inclusive over the kinematics of the unobserved state X , the factorization theorem states that the total inclusive cross-section can be written as:

$$\sigma_{AB \rightarrow f+X} = \sum_{a,b} \int_0^1 dx_a \int_0^1 dx_b f_{a/A}(x_a, \mu_F) f_{b/B}(x_b, \mu_F) \hat{\sigma}_{ab \rightarrow f+X_p}(x_a p_A, x_b p_B, \mu_R, \mu_F) + \mathcal{O}\left(\frac{\Lambda_{\text{QCD}}}{Q}\right)^n, \quad (4.1)$$

This formula represents a convolution of two distinct components, separated by an arbitrary energy scale known as the factorization scale, μ_F . The terms $f_{a/A}(x, \mu_F)$ are the Parton Distribution Functions (PDFs)¹. They represent the probability density of finding a parton of flavor a within hadron A , carrying a longitudinal momentum fraction x of the parent hadron's total momentum. PDFs encode all the long-distance dynamics of the proton. Crucially, PDFs are process-independent quantities. A PDF extracted from Deep Inelastic Scattering (DIS) data at HERA can be used to predict Higgs production at the LHC. This universality is what gives the factorization theorem its immense predictive power. While the PDFs $f_{a/A}(x, \mu_F)$ cannot be calculated perturbatively, their dependence on the factorization scale μ_F is entirely governed by perturbative QCD. This is described by the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) evolution equations [152, 153]:

$$\mu_F^2 \frac{\partial}{\partial \mu_F^2} f_{i/A}(x, \mu_F) = \sum_j \int_x^1 \frac{dz}{z} P_{ij} \left(\frac{x}{z}, \alpha_s(\mu_F) \right) f_{j/A}(z, \mu_F). \quad (4.2)$$

Here, $P_{ij}(x/z)$ are the Altarelli-Parisi splitting functions, calculable in perturbation theory. They represent the probability of a parton j radiating a collinear parton and turning into a parton i . The DGLAP equations allow us to measure PDFs at a low scale μ_0 (via experimental fits) and mathematically evolve them up to the highly energetic scale Q of the LHC collision.

The term $\hat{\sigma}_{ab \rightarrow f+X_p}$ is the partonic cross-section for the constituent partons a and b to produce the final state f plus some additional radiation X_p . Because it takes place at the hard scale $Q \gg \Lambda_{\text{QCD}}$, asymptotic freedom guarantees that $\alpha_s(Q) \ll 1$. Therefore, $\hat{\sigma}$ can be computed as a perturbative series in α_s and α . The differential cross-section $d\sigma$ represents the transition probability per unit incident flux. It connects the quantum mechanical transition amplitude to the experimentally measurable scattering rate. The master formula for the differential cross-section of a $2 \rightarrow n$ process is given by:

$$d\sigma = \frac{1}{\mathcal{F}} |\mathcal{M}|^2 d\Phi_n. \quad (4.3)$$

This equation cleanly factorizes the computation into three distinct components: the kinematic flux factor $\mathcal{F}$, the squared invariant matrix element $|\mathcal{M}|^2$, which is exactly where the multi-loop amplitudes and Feynman integrals discussed in Chapters 2 and 3 are applied, and the Lorentz Invariant Phase Space $d\Phi_n$. Moreover, $\hat{\sigma}$ explicitly depends on the renormalization scale μ_R (arising from the UV regularization of the loop amplitudes) and the factorization scale μ_F .

The final term, $\mathcal{O}(\Lambda_{\text{QCD}}/Q)^n$, represents higher-twist effects (multi-parton interactions, hadronization corrections). For processes where Q is large, these power corrections are strongly suppressed and are typically ignored in high-precision high-energy phenomenology.

The introduction of the factorization scale μ_F is essentially a renormalization procedure applied to infrared, specifically collinear, divergences. When calculating the partonic cross-section $\hat{\sigma}$ beyond Leading Order (LO), one encounters singularities when initial-state partons emit massless gluons collinearly, or when initial-state gluons split collinearly into quark-antiquark pairs. Unlike final-state collinear singularities which cancel for inclusive observables (guaranteed by the KLN theorem), initial-state collinear singularities do not cancel. Instead, these uncanceled logarithmic collinear singularities are systematically factorized and absorbed into the definition of the scale-dependent PDFs. The factorization scale μ_F serves as the arbitrary energy scale at which this separation between long-distance and short-distance physics is performed. Rather than artificially classifying emissions as resolvable or unresolvable based on a kinematic boundary, the proper theoretical framework relies on a rigorous matching procedure.

¹For reviews on PDFs see, for example, Ref. [147–151]

Specifically, the collinear emissions evaluated in the perturbative partonic cross-section are systematically matched to those already absorbed into the PDFs. Through an exact subtraction of these overlapping singular pieces, the collinear divergences are cleanly removed from the partonic calculation, yielding a finite, short-distance partonic cross-section $\hat{\sigma}$. Because μ_F is an unphysical artifact of our calculation method, the true physical hadronic cross-section $\sigma_{ab \rightarrow X}$ must be independent of it:

$$\mu_F^2 \frac{\partial}{\partial \mu_F^2} \hat{\sigma}_{ab \rightarrow X} = 0. \quad (4.4)$$

In practice, because we truncate our perturbative expansion for $\hat{\sigma}$ at a fixed order, a residual μ_F dependence remains. Varying μ_F (typically by factors of 2 around the hard scale Q) is a standard method used to estimate the Missing Higher-Order Uncertainties of the theoretical prediction.

Although originally developed to tackle the non-perturbative complexities of hadron-hadron collisions, the factorization theorem can be naturally generalized to lepton colliders, such as e^+e^- or $\mu^+\mu^-$ machines. This generalization requires substituting the traditional hadronic PDFs with lepton PDFs. The conceptual architecture of the theorem remains identical in both environments: it isolates the short-distance hard scattering amplitude from the long-distance initial-state dynamics, absorbing collinear singularities into distribution functions whose scale dependence is governed by DGLAP-like evolution equations. However, a fundamental difference lies in their physical origin. While proton PDFs encodes the non-perturbative dynamics of color confinement and must be extracted from experimental data, leptons are elementary particles. Consequently, lepton PDFs are purely perturbative quantities. In particular, they account for the resummation of large collinear logarithms, arising from Initial State Radiation (ISR) of photons, which can be rigorously calculated from first principles within quantum electrodynamics and the electroweak theory.

4.2 NLO calculations

To better explain how to perform an higher order calculation, we explicitly go through the computation of NLO QCD corrections to Neutral-Current Drell-Yan (DY), i.e. the production of a lepton pair via an intermediate neutral electroweak gauge boson. A complete introduction to DY will be given in Chapter 5.

At Leading Order (LO), the partonic process consists of quark-antiquark annihilation: $q\bar{q} \rightarrow \ell^+\ell^-$. The corresponding cross-section is integrated over the 2-body final-state phase space, $d\Phi_2$. To achieve NLO accuracy, formally $\mathcal{O}(\alpha_s)$ relative to the Born process, we must evaluate two distinct classes of partonic corrections: the virtual and the real contributions.

The virtual correction arises from the interference between the tree-level Born amplitude and the one-loop amplitude. For DY, this involves a single virtual gluon exchanged between the initial-state quark lines (the vertex correction), as well as quark self-energy diagrams. Because the final state still consists only of the lepton pair, this contribution shares the exact same 2-body phase space as the Born level. However, integrating the loop momentum in exactly 4 dimensions yields divergences in the infrared region (soft and collinear limits). Regulating the theory in $d = 4 - 2\epsilon$ dimensions isolates these divergences as explicit poles of the form $1/\epsilon^2$ and $1/\epsilon$.

The real emission contribution arises from tree-level processes where an additional parton is radiated into the final state. At $\mathcal{O}(\alpha_s)$, two partonic channels open up. Namely,

$$q\bar{q} \rightarrow \ell^+ \ell^- g, \quad (4.5)$$

$$qg \rightarrow \ell^+ \ell^- q. \quad (4.6)$$

These processes are integrated over a 3-body phase space $d\Phi_3$. While the matrix elements for these processes are finite for generic, widely separated kinematics, they become kinematically singular in soft and collinear regions. For instance, the $q\bar{q} \rightarrow \ell^+ \ell^- g$ matrix element diverges when the emitted gluon becomes soft or perfectly collinear with either incoming quark.

The Kinoshita-Lee-Nauenberg theorem guarantees that for any sufficiently inclusive observable, the infrared singularities present in the virtual loop integrals will exactly cancel against the kinematic singularities arising from the integration over the unresolved real-emission phase space. However, it is crucial to emphasize that in hadronic collisions, this automatic cancellation applies strictly to soft singularities and final-state collinear singularities. The initial-state collinear singularities do not cancel between real and virtual corrections. Instead, all infrared divergences are only fully removed after these remaining collinear singularities are factorized and absorbed into the PDFs. If we were calculating the totally inclusive cross-section, we could perform the phase space integrals analytically in d dimensions, observe the exact cancellation of the $1/\epsilon$ poles, and safely take the limit $\epsilon \rightarrow 0$ to obtain a finite result. However, modern precision phenomenology is almost never fully inclusive. Experimental collaborations measure fiducial cross-sections, imposing complex kinematic cuts on the final-state leptons to suppress backgrounds and adapt for detector constraints. The application of these arbitrary experimental cuts, represented mathematically by a non-analytic measurement function, makes the analytical integration of the phase space impossible. We are forced to evaluate the integrals numerically using Monte Carlo methods. This presents a critical bottleneck since we cannot perform numerical integration in 4 dimensions over quantities that are separately divergent.

To solve this, the numerical integration must be restructured such that the virtual and real contributions are made independently finite in 4 dimensions before the Monte Carlo integration takes place. The standard modern approach is the subtraction method [154–156]. The idea is to introduce a local subtraction counterterm $d\hat{\sigma}_S$, which is cleverly constructed to satisfy two requirements. The first one is that it must exactly mimic the singular behavior of the real-emission matrix element $d\hat{\sigma}_R$ in all soft and collinear limits. The second is that it must be analytically integrable over the 1-body phase space of the unresolved radiation in d dimensions. By adding and subtracting this term, the NLO cross-section is reorganized as follows:

$$\hat{\sigma}_{\text{NLO}} = \int_{d\Phi_{n+1}} d\hat{\sigma}_R + \int_{d\Phi_n} d\hat{\sigma}_V = \quad (4.7)$$

$$= \int_{d\Phi_{n+1}} [d\hat{\sigma}_R - d\hat{\sigma}_S + d\hat{\sigma}_S] + \int_{d\Phi_n} d\hat{\sigma}_V = \quad (4.8)$$

$$= \int_{d\Phi_{n+1}} [d\hat{\sigma}_R - d\hat{\sigma}_S] + \int_{d\Phi_n} \left[d\hat{\sigma}_V + \int_{d\Phi_1} d\hat{\sigma}_S \right] \quad (4.9)$$

This structure resolves the numerical bottleneck. In the first integral, the counterterm exactly cancels the divergences of the real emission, rendering the bracketed term completely finite and safely integrable numerically in 4 dimensions. In the second integral, the analytical integration of the counterterm over the unresolved subspace yields explicit $1/\epsilon$ poles. These poles analytically cancel the explicit $1/\epsilon$

poles inherent to the virtual correction. Any remaining initial-state collinear poles are absorbed by the renormalization the bare PDFs.

Crucially, the choice of the subtraction counterterm $d\hat{\sigma}_S$ is not unique. The fundamental requirement of a valid subtraction scheme is just that it matches the real-emission matrix element in all soft and collinear limits. Consequently, one is completely free to add any finite piece to the counterterm. Such an arbitrary modification simply reshuffles finite $\mathcal{O}(\alpha_s)$ contributions between the numerical phase-space integral of the real and the virtual component, leaving the total physical cross-section unchanged. This theoretical freedom has led to the development of several distinct subtraction formalisms, each defining the finite parts and the mapping of the unresolved phase space in different ways. At NLO, two universally adopted, process-independent frameworks dominate the field. Namely, the FKS subtraction proposed by Frixione, Kunszt, and Signer [154], and the dipole subtraction formalism by Catani and Seymour [155, 156].

4.3 NNLO calculations

When evaluating Eq. 4.1 at NNLO even more channels open up. For example, if we consider mixed QCD-EW corrections, i.e. corrections of $\mathcal{O}(\alpha_s\alpha)$ with respect to Born, we have to compute loop corrections for double virtual and real-virtual contributions:

$q\bar{q} \rightarrow \ell^+\ell^-, \gamma\gamma \rightarrow \ell^+\ell^-$	two-loop QCD-EW
$q\bar{q} \rightarrow \ell^+\ell^-g, qg \rightarrow \ell^+\ell^-q$	one-loop EW
$q\bar{q} \rightarrow \ell^+\ell^-\gamma, q\gamma \rightarrow \ell^+\ell^-q, q\gamma \rightarrow \ell^+\ell^-q$	one-loop QCD

and double real processes at tree-level:

$q\bar{q} \rightarrow \ell^+\ell^-g\gamma,$	$q\gamma \rightarrow \ell^+\ell^-qg,$	$g\gamma \rightarrow \ell^+\ell^-q\bar{q},$
$qg \rightarrow \ell^+\ell^-q\gamma,$	$q\bar{q} \rightarrow \ell^+\ell^-q\bar{q},$	$q\bar{q} \rightarrow \ell^+\ell^-q'\bar{q}',$
$qq \rightarrow \ell^+\ell^-qq,$	$qq' \rightarrow \ell^+\ell^-qq',$	$q\bar{q}' \rightarrow \ell^+\ell^-q\bar{q}'.$

The double-real contribution $d\hat{\sigma}_{RR}$ is integrated over the $(2+2)$ -body phase space, $d\Phi_{2+2}$. This term diverges kinematically when up to two partons become simultaneously soft and/or collinear. The Real-Virtual Contribution $d\hat{\sigma}_{RV}$ consists of the interference between one-loop and tree-level amplitudes for the production of $2+1$ partons. It contains explicit $1/\epsilon$ poles from the loop integration (evaluated in $d = 4 - 2\epsilon$ dimensions) and implicit kinematic singularities when the single extra real parton becomes soft or collinear. The Double-Virtual Contribution $d\hat{\sigma}_{VV}$ contains the interference of two-loop amplitudes with tree-level amplitudes, as well as the squared one-loop amplitudes, for the Born kinematics. These multi-loop amplitudes contain deep explicit infrared poles up to $1/\epsilon^4$.

4.3.1 q_T -subtraction

Extending the local, point-by-point subtraction algorithms introduced in previous Section to NNLO is a task of formidable complexity. At NNLO, one must handle double-real emissions where up to two partons can become simultaneously soft and/or collinear in highly intricate, overlapping configurations. To bypass the difficulties of constructing universal local counterterms for the double-real phase space, alternative approaches based on phase-space slicing have been developed. For the calculation of the Drell-Yan process presented in this thesis, we employ the q_T -subtraction formalism [157–159]. Rather

than subtracting singularities point-by-point across the entire phase space, q_T -subtraction is a non-local slicing method that exploits the universal properties of QCD radiation in the limit of small transverse momentum.

The fundamental idea of q_T -subtraction is the use of the transverse momentum of the produced color-singlet final state, denoted as q_T , as an infrared resolution parameter. If we consider the inclusive production of a generic color-singlet system f in a hadronic collision, at Born level we have that the color-singlet is produced via the annihilation of initial-state partons with exactly zero transverse momentum, i.e. $q_T = 0$. Momentum conservation dictates that f can only acquire a non-zero transverse momentum if it recoils against additional real partonic radiation emitted into the final state. Therefore, the condition $q_T > 0$ strictly bounds the kinematics away from the purely virtual Born-level configuration. If we introduce a small theoretical cutoff, q_T^{cut} , the phase space naturally partitions into two distinct regimes:

- a resolved regime with $q_T > q_T^{\text{cut}}$ in which the final state must contain at least one resolved, hard parton. The NNLO calculation for the inclusive production of f is then kinematically reduced to an NLO calculation for f +jet;
- an unresolved regime in which $q_T < q_T^{\text{cut}}$. In this case, the system is accompanied by purely soft and/or collinear radiation.

By combining these two regimes, the q_T -subtraction formalism expresses the fully differential NNLO cross-section for the production of a color-singlet state f as:

$$d\hat{\sigma}_{\text{NNLO}}^f = \mathcal{H}_{\text{NNLO}}^f \otimes d\hat{\sigma}_{\text{LO}}^f + \left[d\hat{\sigma}_{\text{NLO}}^{f+\text{jet}} - d\hat{\sigma}_{\text{NNLO}}^{\text{CT}} \right]. \quad (4.10)$$

This master formula breaks the complex NNLO calculation into three manageable, automatable components. The real emission term, $d\hat{\sigma}_{\text{NLO}}^{f+\text{jet}}$, represents the NLO cross-section for the production of f accompanied by at least one hard jet. Because this is an NLO quantity, it can be evaluated using standard local subtraction techniques without requiring any new theoretical machinery. In the limit $q_T \rightarrow 0$ this real contribution is divergent, since the recoiling radiation becomes soft and/or collinear to the initial-state partons. Such divergence is cancelled by the counterterm $d\hat{\sigma}_{\text{NNLO}}^{\text{CT}}$. More precisely, $d\hat{\sigma}_{\text{NNLO}}^{\text{CT}}$ guarantees a non-local cancellation at the level of the inclusive q_T distribution, i.e. at any fixed value of $\vec{q}_T$ but only upon integration over all other kinematic degrees of freedom. Finally, the hard function $\mathcal{H}_{\text{NNLO}}^f$ collects the purely virtual two-loop corrections to the Born process, along with the finite remainders left over after the analytic cancellation of the infrared poles between the two-loop integrals and the subtraction terms integrated down to $q_T = 0$. The calculation of the hard function will be the focus of Chapter 5. In that Chapter we will also present the integrated counter-terms which makes the hard function finite. The quantity in the square brackets in Eq. 4.10 represents the difference between the real-emission cross-section and the subtraction counterterm. The counterterm is designed to reproduce the exact singular behavior of the real emission, both components are individually divergent, however their difference is finite in the $q_T \rightarrow 0$ limit. Crucially, due to the non-local nature of the q_T -subtraction method, this exact cancellation cannot be evaluated directly at $q_T = 0$ point-by-point in a four-dimensional Monte Carlo integration without encountering severe numerical instabilities. This is precisely why a finite cutoff parameter, $r_{\text{cut}} = q_T^{\text{cut}}/Q$, is required in practical applications. The numerical integration of the quantity in the square brackets is instead restricted to the kinematically well-defined region $q_T > q_T^{\text{cut}}$. However, the presence of this cutoff introduces unphysical power corrections in r_{cut} to the final cross-section. To obtain a cutoff-independent prediction, the numerical phase-space integration must be performed for several progressively smaller values of r_{cut} . The physical result is then extracted by extrapolating the dependence of the cross-section to the $r_{\text{cut}} \rightarrow 0$ limit.

4.4 Automatization of NNLO calculations

In the previous sections, we investigated the conceptual and mathematical anatomy of an NNLO cross-section. We established that predicting a physical observable at NNLO accuracy requires the computation of double-virtual, real-virtual, and double-real matrix elements, followed by a highly non-trivial extraction and cancellation of infrared singularities.

For NLO calculations, this entire pipeline, from the generation of Feynman diagrams to the final phase-space integration, has been fully automated in general-purpose Monte Carlo event generators like MadGraph5_aMCNLO [20, 160]. The user simply specifies a process, and the software handles the rest, relying on one-loop amplitude providers and universal subtraction schemes like Catani-Seymour dipole subtraction or the FKS method.

Automating NNLO calculations in a similar manner is the current frontier of high-energy phenomenology.

4.4.1 Two-loop amplitudes

Here we present OCEANN [6], a new complete workflow for the evaluation of two-loop corrections. The computational framework OCEANN is a collection of codes and packages that provides, to an experienced user, all the necessary tools to face the challenges of the calculation of a multi-loop electroweak amplitude in a consistent way. The main steps of these computations are: the calculation of the squared matrix element which contains all the interference terms between the Feynman diagrams that contribute to the process; the evaluation of the Feynman loop integrals present in the amplitude; the cancellation of the UV and IR divergences via the renormalisation and subtraction procedures and, finally preparation of numerical grids. Because the exact evaluation of multi-loop amplitudes is highly complex and computationally expensive, calculating them on-the-fly at every single phase-space point sampled by a Monte-Carlo generator would require an unfeasible amount of time. Numerical grids solve this bottleneck. They are essentially pre-computed datasets that map the value of the finite amplitude across a discrete, multidimensional mesh of phase-space points. Once these grids are generated upstream, Monte-Carlo integrators can rely on fast interpolation algorithms to evaluate the amplitude at any arbitrary kinematic configuration. This approach decouples the slow amplitude evaluation from the phase-space integration, drastically reducing the total runtime while preserving high theoretical precision.

The evaluation of the squared matrix elements is implemented, within OCEANN, in the Mathematica package ABISS (Amplitudes Builder, Interference Solver and Simplifier). ABISS expresses the final result as a combination of MIs, each multiplied by a rational function of the kinematical invariants and masses of the process, with a workflow which can be described schematically by the following points:

1. The generation of the Feynman diagrams of the chosen scattering process is performed with FEYNARTS [161], which supports several model files for the SM and its extensions.
2. The Feynman diagrams with the insertion of UV counterterms are generated with a custom model file for FEYNARTS which provides the counterterms in symbolical form. Their numerical values, as a Laurent series in the dimensional regularization parameter ϵ , can be computed together with all the other diagrams, or they can be read from any external counterterm library.

3. The evaluation of the interference between an amplitude and its Hermitian conjugate², including the sum over all the external polarizations, is performed by dedicated routines in ABISS. The result is a list of expressions, Lorentz scalars, including Feynman loop integrals and coefficients, functions of kinematical invariants and masses.
4. All the Feynman integrals are classified in terms of integral families. The user can propose a preferred choice of integral families, in terms of which ABISS tries to write all the Feynman integrals present in the computation³. If the assignment of a given Feynman integral to one of the available families is not possible, ABISS automatically completes the tentative set, using the information of the new Feynman integral to introduce the new family. At the end of this classification step, the initial information relevant for the reduction to a minimal set of so called MIs is available and automatically organized in a group of dedicated directories.
5. The output of the previous point is ready for the external run of any code able to generate all the relations useful to express a generic element of an integral family as a linear combination of MIs. Such relations include IBP and Lorentz identities. Within OCEANN, we consider as reference code for this task KIRA [112–114, 120]. The outcome of the run can be imported back to the directories already set-up by ABISS.
6. With the help of the reduction rules computed in the previous point, ABISS is able to organize the squared matrix element as a sum over the N_{MI} MIs $\mathcal{I}_j$, with appropriate coefficients c_j . For example, in the case of the two-loop contributions,

$$\langle \mathcal{M}^{(0)} | \mathcal{M}^{(0,2)} \rangle = \sum_{j=1}^{N_{MI}} c_j \mathcal{I}_j . \quad (4.11)$$

Both coefficients and integrals depend on the kinematical invariants and on the masses of the particles of the model.

The evaluation of the MIs which have been identified in the first stage is handled within OCEANN with the public Mathematica package SEASYDE [2, 5], following a simple workflow:

1. For each integral family, we compute the systems of differential equations satisfied by its MIs, with respect to the kinematical invariants. For this task, several tools are publicly available, like e.g. LITERED [115] or REDUZE [117].
2. We compute the boundary condition values for each MI with the code AMFLOW [130].
3. We solve the system of differential equations with SEASYDE.

SEASYDE solves the system of differential equations with the series expansion procedure [106, 132]. The solution is studied in the complex plane, with an original algorithm for the analytical continuation of the solutions. The efficient evaluation of a series of points in one single run is relevant for the preparation of numerical grids which sample the whole phase space of the process. The numerical results have arbitrary precision, which can be selected by the user and is achieved with an appropriate number of terms in the power series used to represent the solution.

²We can in full generality consider an amplitude and a projector.

³Different routings of the loop momenta and trivial symmetry relations are tested, to match the momenta with those of the definition of the integral family.

Once the interference terms have been computed and the Feynman integrals evaluated, then the last steps of the calculation are handled by internal routines within OCEANN: the combination of the diagrammatic and counterterm contributions, to achieve the UV finiteness of the amplitude; the subtraction of the IR poles from the UV renormalised interference terms. The sequences of points, covering the whole phase space, computed by SEASYDEFOR for the MIs and within ABISS for their rational coefficients, express the finite contribution of the virtual corrections, relevant for phenomenological studies. One possible strategy is the preparation of grids which allow the evaluation of the correction at any phase-space point via interpolation, with excellent accuracy.

This approach has been applied, within a preliminary version of the OCEANN framework, for example, in the context of the mixed QCDxEW corrections in [1, 3], providing in particular the validation of this new semi-analytical approach, to handle the cancellation of UV and IR divergences, represented by poles in dimensional regularization.

The sequence of computational steps described above, taking place within the OCEANN framework, is fully general and applies to the evaluation of the purely virtual corrections of an arbitrary scattering process at a given perturbative order.

4.4.2 Phase-space integration

With the automated generation of two-loop amplitudes addressed, the final critical step in producing a physical prediction is the integration of these matrix elements over the multi-particle phase space. At high-energy particle colliders, the inclusive cross-section takes the form of Eq. 4.1.

However, in a realistic collider environment, detectors have finite acceptances, e.g. particles escaping down the beam pipe are unmeasured, and large backgrounds necessitate strict fiducial cuts, such as requiring a minimum transverse momentum p_T or restriction on the pseudo-rapidity η of final-state leptons and jets. These experimental realities are encoded in a measurement function $\Theta_{\text{cuts}}(p_1, \dots, p_n)$, which is added to the integrand of Eq. 4.1. This measurement function acts as a complex, non-analytic, multi-dimensional step function. Because these arbitrary geometric boundaries make analytic integration impossible, the field must rely entirely on Monte Carlo integration. Monte Carlo generators evaluate the phase space by randomly sampling multi-particle kinematic configurations (events). For each event, the program checks if the kinematics pass the Θ_{cuts} filter; if so, the event is kept and weighted by the exact value of the squared matrix element. Furthermore, the dimensionality of the phase space at NNLO is exceptionally high; a $2 \rightarrow 4$ process evaluated at the double-real emission level requires integrating over a computationally heavy multi-particle phase space. This numerical approach, hence, is the only mathematically viable way to translate d -dimensional QFT into highly differential predictions that can be directly plotted against LHC data.

While general-purpose event generators, like MadGraph5_aMCNLO, have completely automated this integration process at NLO, extending this automation to NNLO is a formidable challenge due to the delicate numerical cancellations required between real emissions and IR subtraction counterterms in deep unresolved limits. Despite these challenges, highly successful specialized frameworks have emerged. The state-of-the-art example for computing NNLO QCD and NNLO QCD-EW and NLO EW corrections is the MATRIX (Munich Automates q_T subtraction and Resummation to Integrate Cross Sections) framework [162]. MATRIX effectively achieves a "push-button" automation for a wide class of processes, specifically the production of massive, color-singlet final states (such as Drell-Yan, Higgs, and diboson production).

MATRIX automates the phase-space integration by leveraging the q_T -subtraction slicing formalism. It coordinates the calculation by dynamically partitioning the phase space: it utilizes the Munich

Monte Carlo integrator to handle the highly structured multi-channel integration of the resolved real-emission process alongside automated dipole subtraction, while simultaneously integrating the heavily peaked analytical $q_T \rightarrow 0$ counterterms. By interfacing with automated one-loop providers like OPEN-LOOPS [163], RECOLA [164] and various specialized two-loop amplitude libraries, such as the one provided by the OCEANN framework, MATRIX allows phenomenologists to define arbitrary experimental cuts and extract fully differential NNLO distributions in a completely automatic way.

Two-loop amplitudes for Drell-Yan

“If you do things properly, climbing is not dangerous. It is just a matter of extreme precision.”

Ueli Steck, the "Swiss machine", *alpinist*.

In this Chapter we present our calculation for the two-loop mixed corrections to NC- [1] and CC-DY [3], the two-loop QED ones to NC-DY [6] and we will give some prospects towards full two-loop EW calculation.

The Drell-Yan (DY) processes [165] are standard candles at colliders, such as the Large Hadron Collider (LHC) in Geneva, Switzerland. They consist in the hadroproduction of a lepton pair mediated by a vector boson. When the mediating boson is electrically neutral, the process is referred to as Neutral Current DY (NC-DY); when it carries electric charge, it is known as Charged Current DY (CC-DY).

At Leading Order, the DY process proceeds via the s-channel annihilation of a quark and an antiquark from the colliding hadrons. The two partonic processes are:

$$\begin{aligned} q\bar{q} &\rightarrow \ell^+ \ell^- \\ q\bar{q}' &\rightarrow \ell^\pm \nu \end{aligned}$$

To describe the kinematics of the DY process, we define the four-momenta of the incoming hadrons as P_1 and P_2 . The center-of-mass energy squared of the hadronic system is given by $s = (P_1 + P_2)^2$. The annihilating quark and antiquark carry momentum fractions x_1 and x_2 of their respective parent hadrons, such that their four-momenta are $p_1 = x_1 P_1$ and $p_2 = x_2 P_2$. The invariant mass squared of the partonic system, $\hat{s}$, is directly related to the hadronic center-of-mass energy:

$$\hat{s} = (p_1 + p_2)^2 = x_1 x_2 s, \quad (5.1)$$

and is exactly the invariant mass of the off-shell vector-boson, $\hat{s} = m_V^2$. The latter, in case of the NC-DY is fully reconstructible from the four momenta of the two leptons:

$$m_V^2 = (p_{\ell_1} + p_{\ell_2})^2. \quad (5.2)$$

Another experimentally accessible variable is the boson rapidity, y_V . It represents the velocity component along the beam axis of produced bosons relative to the speed of light. It measures how far forward, or backward, a boson is produced relative to the hadronic collision center ($y_V = 0$). This is defined as:

$$y_V = \frac{1}{2} \log \left(\frac{E_V + p_{z,V}}{E_V - p_{z,V}} \right) = \frac{1}{2} \log \left(\frac{x_1}{x_2} \right). \quad (5.3)$$




Figure 5.1: Schematic picture of the Neutral Current Drell-Yan process.

By combining the equations for m_V^2 and y_V , the initial partonic momentum fractions can be mapped directly to the experimentally observable macroscopic variables:

$$x_{1,2} = \frac{m_V}{\sqrt{s}} e^{\pm y_V}. \quad (5.4)$$

The transverse momentum of the boson is the magnitude of the boson's momentum projected onto the plane orthogonal to the beam axis.

$$p_{T,V} = \sqrt{p_{x,V}^2 + p_{y,V}^2}. \quad (5.5)$$

At strictly LO, the vector boson is produced with zero transverse momentum, $p_{T,V} = 0$, because the initial partons are purely collinear. In experimental reality, a non-zero $p_{T,V}$ is observed due to intrinsic transverse momentum of partons within the proton and, more dominantly, hard radiation at higher orders.

The decay of the vector boson yields leptons that are directly measured by the detector. The kinematics of these final-state particles are typically described by means of its transverse momentum and pseudorapidity:

$$p_{T,\ell} = \sqrt{p_{x,\ell}^2 + p_{y,\ell}^2},$$

$$\eta_\ell = -\log \left[\tan \left(\frac{\theta}{2} \right) \right]. \quad (5.6)$$

where θ is the polar angle relative to the beam axis. Finally, in CC processes, the neutrino escapes the detector without interacting. Its presence is inferred through missing transverse momentum. Because the momentum of the neutrino cannot be measured, the true invariant mass of the W boson cannot be reconstructed. Instead, the transverse mass $m_{T,V}$ is used, combining only the variables in the transverse plane:

$$m_{T,V} = \sqrt{2p_{T,\ell} p_T^{\text{miss}} (1 - \cos \Delta\varphi)} \quad (5.7)$$

where $\Delta\varphi$ is the azimuthal angle between the charged lepton and the missing transverse momentum vector. These quantities are fundamental in phenomenological studies, and they will be further investigated in Chapter 6.

The importance of DY in collider physics comes to the fact that the production rates are large, the experimental signatures are clean, given the presence of at least one lepton with large transverse momentum in the final state, and there is no color flow between initial and final state. For this reason, they provide the perfect environment for performing precision studies of the gauge sector of the SM. In particular, the NC-DY is important for the determination of the Z -boson mass and the effective weak mixing angle,

while the CC-DY is important for the determination of the W -boson mass. These three EW parameters are known at the moment with small relative errors, 0.01% for the boson masses [166–168] and 0.1% for the effective weak mixing angle [169, 170], originating from a fitting procedure between measured kinematical distributions and generated theoretical templates [171–175]. The DY processes are also crucial for the extractions of parton distribution functions (PDFs) [176] and the strong coupling constant α_s [177, 178]. It is, hence, of extreme importance to minimize the theoretical uncertainties on physical predictions in order to make the comparison between theory and experiment more significant. The direct search for signals of New Physics (NP) is carried on, in the DY process, with the precise measurement of the tail of the kinematical distributions. A deviation could then be interpreted as the distortion due to a new resonance or, in general, as the effect of a new higher-dimensional operator whose growth with the energy shows up in the tails. In order to appreciate the possibility of such a deviation, the most precise SM prediction is needed, with the smallest possible residual theoretical uncertainty. The inclusion of higher-order corrections makes the theoretical predictions more precise, but it requires to overcome challenging mathematical and computational problems.

The DY process was one of the first hadronic reactions for which radiative corrections in the strong and EW couplings α_s and α were computed. Classic calculations of the NLO [179] and NNLO [180, 181] corrections to the total cross section in QCD were followed by (fully) differential NNLO computations including the leptonic decay of the vector boson [182–185]. The complete EW corrections for W production have been computed in Refs. [186–190], and for Z production in Refs. [191–195]. Recently, the next-to-next-to-next-to-leading-order (N^3 LO) QCD radiative calculations for the inclusive production of a virtual photon [196, 197] and of a W boson [198] have been completed, and also computations of fiducial cross sections at this order were performed [199–202].

The mixed QCD-QED corrections to the inclusive production of an on-shell Z boson were obtained in Ref. [203] through an Abelianisation procedure from the NNLO QCD results [180, 181], and later extended in Ref. [204] to the fully differential level for off-shell Z boson production and decay into a pair of neutrinos, thereby avoiding final-state radiation. A similar calculation was carried out in Ref. [205] in an on-shell approximation for the Z boson, but including the factorised NLO QCD corrections to Z production and the NLO QED corrections to the leptonic Z decay. Complete $\mathcal{O}(\alpha\alpha_s)$ computations for the production of on-shell Z and W bosons have been presented in Refs. [206–211]. Beyond the on-shell approximation, important results have been obtained in the *pole* approximation (see e.g. Ref. [27] and references therein). In this approximation the cross section around the W or Z resonance is systematically expanded so as to split the radiative corrections into well-defined, gauge-invariant contributions. Such method has been used in Refs. [212, 213] to evaluate the dominant part of the mixed QCD–EW corrections in the resonant region. The calculation in this approximation was recently completed in Ref. [214].

A first step beyond the pole approximation has been carried out in Ref. [73], where results for the $\mathcal{O}(n_F\alpha_s\alpha)$ contributions to the DY cross section were presented. In Ref. [215] some of us presented a computation of the mixed QCD–EW corrections to the charged-current process $pp \rightarrow \ell\nu_\ell + X$, where all contributions are evaluated exactly except for the finite part of the two-loop amplitude, which is evaluated in the pole approximation. Recently, the exact evaluation of the two-loop amplitude for the QCD–EW corrections to this process has been presented in Ref. [3].

The complete computation of the mixed QCD–EW corrections to the neutral-current process $pp \rightarrow \ell^+\ell^- + X$ has been reported in Ref. [216] considering massive leptons, and in Ref. [217] for massless leptons. The calculation of Ref. [216] is based on our exact two-loop amplitudes Ref. [1], evaluated through the reduction to the master integrals of Ref. [138] and a semi-analytical implementation of the

differential-equations method for their calculation [2, 106, 132]. The computation of Ref. [217] is based on the two-loop amplitudes presented in Ref. [218], expressed in terms of the master integrals evaluated in analytic polylogarithmic form [139, 140].

The role of the mixed NNLO QCD–EW corrections is central for a correct interpretation of high-precision DY data [219–221]. According to the factorization theorem, we can write the hadron level cross section as a convolution of the proton PDFs with a partonic cross section, and both factors feature QCD and EW radiative effects. Initial-state QED and mixed QCD–QED collinear singularities can be factorized from the partonic cross section and reabsorbed in the definition of the physical proton PDF [222, 223]. The coupled DGLAP evolution, with QCD and QED splitting kernels, yields a non-trivial interplay which goes beyond the simple single-emission probability, with mixed higher-order effects automatically included [224–226]. The presence of a photon density in the proton opens the possibility of new partonic scattering channels, necessary to complement the cross sections initiated by quarks and gluons. The relative contribution of the various channels changes, as a function of the invariant mass of the final state, and only an NNLO QCD–EW calculation can consistently account for their correct balance [227–229]. The availability of an exact NNLO QCD–EW calculation throughout the whole invariant-mass range, including both the Z resonance and the large invariant-mass tail, is important to predict the shape of the kinematical distributions, and, in turn, to establish in a statistically significant way any discrepancy from the data. In a distinct and complementary perspective, the increased precision of the partonic results may help putting tighter constraints on the proton parameterization at large partonic x [230]. Moreover, the knowledge of the NNLO QCD–EW corrections potentially allows us to improve over the approximations used in shower Monte Carlo programs (see e.g. Refs. [231, 232], Sec. IV.1 of Ref. [233], and the benchmarking study of Ref. [234]), which include only partial subsets of factorizable mixed QCD–EW corrections, and to reduce the remaining theoretical uncertainties.

5.1 Two-loop mixed corrections to NC-DY

5.1.1 Framework of the calculation

We consider two-loop mixed QCD–EW corrections to lepton pair production in the quark annihilation channel, as given by

$$q(p_1) + \bar{q}(p_2) \rightarrow l^-(p_3) + l^+(p_4). \quad (5.8)$$

The bare amplitude¹ for this partonic process admits a double perturbative expansion in the two coupling constants

$$|\hat{\mathcal{M}}\rangle = 4\pi\hat{\alpha} \sum_{i,j} \left(\frac{\hat{\alpha}_s}{4\pi}\right)^i \left(\frac{\hat{\alpha}}{4\pi}\right)^j |\hat{\mathcal{M}}^{(i,j)}\rangle. \quad (5.9)$$

Here we present the calculation of the following interference terms:

$$\langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(1,0)} \rangle, \quad \langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(0,1)} \rangle, \quad \langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(1,1)} \rangle, \quad (5.10)$$

which contribute at $\mathcal{O}(\alpha\alpha_s)$ to the unpolarised squared matrix element of the process in Eq. 5.8. It is worth noting that, although the first two interference terms formally correspond to the purely one-

¹The *hat* denotes the bare quantities.

loop QCD $\mathcal{O}(\alpha_s)$ and one-loop EW $\mathcal{O}(\alpha)$ corrections, they are strictly required to construct the fully UV-renormalised and IR-subtracted finite contribution at the mixed $\mathcal{O}(\alpha\alpha_s)$ order (see Eq. 5.43).

Virtual corrections are affected by singularities both of UV and IR type. The renormalisation procedure allows us to cancel the UV divergences and to express the amplitude in terms of renormalised parameters, which are in turn related to measurable quantities. The presence of IR divergences in the virtual corrections, on the other hand, is the counterpart of those which appear upon phase-space integration in the processes with the emission of additional massless partons, in our case a photon and/or a gluon. The universality of the IR divergent structure of the scattering amplitudes arising from soft and/or collinear radiation has been widely discussed in the literature and allows us to build a subtraction term, independently of the details of the full two-loop calculation. We exploit this property to check the consistency of our computation, in particular for the prescription related to the Dirac γ_5 matrix. To regularise both UV and IR divergences, we work in $d = 4 - 2\epsilon$ space-time dimensions. The renormalised amplitude can thus be written as a Laurent expansion in ϵ . At one-loop level we have:

$$\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,1)} \rangle = \sum_{i=-2}^2 \epsilon^i C_i^{(0,1)}(s, t, m_W, m_Z, m_\ell), \quad (5.11)$$

$$\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,0)} \rangle = \sum_{i=-2}^2 \epsilon^i C_i^{(1,0)}(s, t, m_W, m_Z, m_\ell), \quad (5.12)$$

while at two-loop we have:

$$\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1)} \rangle = \sum_{i=-4}^0 \epsilon^i C_i^{(1,1)}(s, t, m_W, m_Z, m_\ell). \quad (5.13)$$

In Eq. 5.13 we have retained only terms up to ϵ^0 , as we are only interested in the finite cotributions at NNLO. However, this series contains higher powers of ϵ which will be relevant for higher-order calculations. In these expressions m_ℓ is the lepton mass, and μ_W, μ_Z are the complex masses of the W and Z boson defined as:

$$\mu_V^2 = m_V^2 - im_V \Gamma_V. \quad (5.14)$$

The mass and decay width m_V, Γ_V are real parameters and are the pole quantities, i.e. correspond to the position of the pole in the complex plane of the gauge boson propagator. We also introduce the Mandelstam variables, defined as follows:

$$s = (p_1 + p_2)^2, t = (p_1 - p_3)^2, u = (p_2 - p_3)^2 \text{ with } s + t + u = 2m_\ell^2. \quad (5.15)$$

The on-shell conditions of the external particles are given by

$$p_1^2 = p_2^2 = 0; p_3^2 = p_4^2 = m_\ell^2; \quad (5.16)$$

The coefficients C_i in Eq. 5.13 constitute the IR-unsubtracted two-loop interference terms, while those in Eqs. 5.11-5.12, computed at one-loop, are necessary to prepare the subtraction term. The IR-subtracted expressions that can be obtained by their combination constitutes our main result.

5.1.2 Ultraviolet renormalisation

The renormalisation at $\mathcal{O}(\alpha\alpha_s)$ of the NC DY process has already been discussed in detail in ref. [73]. We report here the basic steps that we have implemented to obtain the complete two-loop renormalised amplitude.

Charge renormalisation

The bare gauge couplings g_0, g'_0 and the Higgs doublet vacuum expectation value v_0 are expressed in terms of their renormalised counterparts g, g', v via the introduction of appropriate counterterms. The relation of g, g', v to a set of three measurable quantities, like for instance G_μ, m_W, m_Z (with G_μ the Fermi constant) or α, m_W, m_Z (with α the fine structure constant)², eventually allows the numerical evaluation of the amplitude. We define for convenience two additional bare quantities: the sine squared of the on-shell weak mixing angle, which we abbreviate as $s_{W0}^2 = \sin^2 \theta_{W0} = 1 - \frac{\mu_{W0}^2}{\mu_{Z0}^2}$, $c_{W0}^2 = 1 - s_{W0}^2$, and the electric charge $e_0 = g_0 s_{W0}$, but we stress that only three of the above parameters are independent. We introduce the relation between bare and renormalised input parameters

$$\mu_{W0}^2 = \mu_W^2 + \delta\mu_W^2, \quad \mu_{Z0}^2 = \mu_Z^2 + \delta\mu_Z^2, \quad e_0 = e + \delta e, \quad (5.17)$$

$$\frac{\delta s_W^2}{s_W^2} = \frac{c_W^2}{s_W^2} \left(\frac{\delta\mu_Z^2}{\mu_Z^2} - \frac{\delta\mu_W^2}{\mu_W^2} \right). \quad (5.18)$$

The on-shell electric charge counterterm δe at $\mathcal{O}(\alpha^2)$ and $\mathcal{O}(\alpha\alpha_s)$ has been discussed in ref. [79], from the study of the Thomson scattering. The mass counterterms $\delta\mu_W^2, \delta\mu_Z^2$ in the complex mass scheme [98] have been presented in ref. [73]. In terms of the transverse part of the unrenormalised VV gauge boson self-energies, they are defined as follows:

$$\delta\mu_V^2 = \Sigma_T^{VV}(\mu_V^2), \quad (5.19)$$

at the pole in the complex plane $q^2 = \mu_V^2$ of the gauge boson propagator, with $V = W, Z$. From the study of the muon-decay amplitude, we derive the following relation

$$\frac{G_\mu}{\sqrt{2}} = \frac{\pi\alpha}{2\mu_W^2 s_W^2} (1 + \Delta r). \quad (5.20)$$

The finite correction Δr was introduced, with real gauge boson masses, in ref. [91] and its $\mathcal{O}(\alpha\alpha_s)$ corrections were presented in ref. [236, 237]. We evaluate it here with complex-valued masses.

We consider now the bare couplings which appear at tree-level in the interaction of the photon and Z boson with fermions. By means of the relations in Eq. 5.17 we obtain the UV divergent correction factor δg_Z^α and $\delta g_Z^{G_\mu}$. The first one contributes to the charge renormalisation of the $Zf\bar{f}$ vertex in the (α, m_W, m_Z) input scheme

$$\frac{g_0}{c_{W0}} = \frac{e}{c_W s_W} \left[1 + \frac{1}{2} \left(2 \frac{\delta e}{e} + \frac{s_W^2 - c_W^2}{c_W^2} \frac{\delta s_W^2}{s_W^2} \right) \right] \equiv \frac{\sqrt{4\pi\alpha}}{c_W s_W} (1 + \delta g_Z^\alpha), \quad (5.21)$$

while the second is relevant in the (G_μ, m_W, m_Z) input scheme

$$\frac{g_0}{c_{W0}} = \sqrt{4\sqrt{2}G_\mu\mu_Z^2} \left[1 - \frac{1}{2}\Delta r + \frac{1}{2} \left(2 \frac{\delta e}{e} + \frac{s_W^2 - c_W^2}{c_W^2} \frac{\delta s_W^2}{s_W^2} \right) \right] \equiv \sqrt{4\sqrt{2}G_\mu\mu_Z^2} (1 + \delta g_Z^{G_\mu}) \quad (5.22)$$

²An alternative input scheme, relevant for the direct determination of the weak mixing angle, has been presented in ref. [235]

Working out the explicit expression of $\delta g_Z^{G_\mu}$, the dependence on the electric charge counterterm cancels out. Analogously, in the case of the $\gamma f \bar{f}$ vertex, the electric charge renormalisation is given by

$$g_0 s_{W0} = e_0 = e + \delta e \equiv \sqrt{4\pi\alpha} (1 + \delta g_A^\alpha) \quad (5.23)$$

in the (α, m_W, m_Z) scheme and by

$$g_0 s_{W0} = \sqrt{4\sqrt{2}G_\mu m_W^2 s_W^2} \left[1 + \frac{1}{2} (-\Delta r + 2\frac{\delta e}{e}) \right] \equiv \sqrt{4\sqrt{2}G_\mu m_W^2 s_W^2} (1 + \delta g_A^{G_\mu}) \quad (5.24)$$

in the (G_μ, m_W, m_Z) scheme. The counterterm contributions to the renormalised amplitude are obtained by replacing the bare couplings in the lower order amplitudes with the expressions presented in eqs. (5.21-5.24) and expanding $\delta g_{A,Z}$ up to the relevant perturbative order. We show in the next Section how $\delta g_{A,Z}$ enter in the renormalisation of the gauge boson propagators.

Renormalisation of the gauge boson propagators

The renormalised gauge boson self-energies are obtained, at $\mathcal{O}(\alpha)$, by combining the unrenormalised self-energy expressions with the mass and wave function counterterms. In the full calculation, we never introduce wave function counterterms on the internal lines, because they would systematically cancel against a corresponding factor stemming from the definition of the renormalised vertices. We exploit instead the relation in the SM between the wave function and charge counterterms [27] and we directly use the latter to define the renormalised self-energies. We obtain:

$$\Sigma_{R,T}^{AA}(q^2) = \Sigma_T^{AA}(q^2) + 2q^2 \delta g_A, \quad (5.25)$$

$$\Sigma_{R,T}^{ZZ}(q^2) = \Sigma_T^{ZZ}(q^2) - \delta\mu_Z^2 + 2(q^2 - \mu_Z^2) \delta g_Z, \quad (5.26)$$

$$\Sigma_{R,T}^{AZ}(q^2) = \Sigma_T^{AZ}(q^2) - q^2 \frac{\delta s_W^2}{s_W c_W}, \quad (5.27)$$

$$\Sigma_{R,T}^{ZA}(q^2) = \Sigma_T^{ZA}(q^2) - q^2 \frac{\delta s_W^2}{s_W c_W}, \quad (5.28)$$

where $\Sigma_T^{VV'}$ and $\Sigma_{R,T}^{VV'}$ are the transverse part of the bare and renormalised VV' vector boson self-energy. The charge counterterms have been defined in Eqs. (5.21-5.24). At $\mathcal{O}(\alpha_s)$ the structure of these contributions does not change: the corrections to the gauge boson self-energies stem from a quark loop with one internal gluon exchange and, in addition, from the $\mathcal{O}(\alpha_s)$ mass renormalisation of the quark lines in the one-loop self-energies.

The expression of the two-loop Feynman integrals required for the evaluation of the $\mathcal{O}(\alpha_s)$ correction to the gauge boson propagators and all the needed counterterms can be found in refs. [73, 236, 237].

Checks

The physical scattering amplitude does not depend on the gauge choice used to quantise the theory. In our computation we use the background field gauge (BFG) [48], which restores some QED-like Ward identities in the full EW SM and guarantees in turn some properties of the Green's function which appear in the calculation. For instance, the UV finiteness of the vertex corrections, summed with the external wave function of the fermions entering in that vertex, is an important property which can provide an additional check of the calculation. This check has been performed explicitly at the one-loop level, based on the knowledge of the UV poles of the one-loop Feynman integrals. At two-loop level, we consider the UV finiteness of the box corrections, we explicitly renormalise the tree-level propagators, and we assume

that the cancellation of the UV poles between vertex and external wave function corrections takes place. Since the propagator corrections are IR finite, we conclude that the sum of vertex, box, and external wave function corrections contains only IR poles. We have explicitly checked that the coefficients of these poles exactly match the ones predicted by the universal structure of the IR divergences in gauge theories [238]. This result hints in favour of the validity of the Ward identity also at two-loop level.

5.1.3 Generation of the full amplitude and evaluation of the interference terms

The amplitude at $\mathcal{O}(\alpha\alpha_s)$ receives contributions from different kinds of Feynman diagrams, for which we depict³ in Figures 5.2 and 5.3 a few representative examples. In the initial state we have two-loop vertex corrections (Figure 5.2-a), which are combined with the two-loop external quark wave function corrections (Figure 5.2-b) and with the one-loop external quark wave function correction with initial-state QCD vertex correction (Figure 5.3-f), yielding an UV-finite, but still IR-divergent, result. The two-loop gauge boson self-energies, together with the corresponding two-loop mass counterterms and with the charge renormalisation constants of the initial and final state vertices are shown in Figures (5.2-c,d,e) and yield an UV- and IR-finite contribution. An example of two-loop boxes with the exchange of neutral or charged EW bosons is given in Figure 5.2-f. At $\mathcal{O}(\alpha\alpha_s)$ all the factorizable contributions




Figure 5.2: Sample Feynman diagrams of two-loop corrections and associated two-loop counterterms.




Figure 5.3: Sample Feynman diagrams of factorisable corrections, including one-loop counterterm corrections.

include an initial-state QCD vertex correction; the second factor can be, alternatively: the final-state EW vertex corrections, the external lepton wave function correction, the one-loop self-energy corrections, the one-loop mass and charge renormalization counterterms, and the one-loop external quark wave function correction. They are schematically represented in Figures 5.3(a)–(f) respectively. Also in this case the

³The straight line with an arrow represents a fermion line (blue: massless quark, red: massive leptons, green: neutrinos), the wavy lines represent EW gauge bosons and the spiral coils represent gluons.

BFG properties allow to identify UV-finite combinations of Feynman diagrams. In the interference term $\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1)} \rangle$ we recognise different subsets associated to different combinations of the EW bosons exchanged in the loops. We consider the following cases: $\gamma\gamma$, γZ , ZZ , WW ; in the neutral cases there are either one boson in a loop and one in a tree-level line or both bosons in the loops, as depicted in a few representative cases in Figure 5.4 a-c for the γZ subset; in the charged case we find one or two W s always in the loop lines Figure 5.4 d-f. The calculation of the bare matrix elements up to two-loop




Figure 5.4: Sample Feynman diagrams contributing to the γZ and to the WW subsets.

follows a general procedure. For the generation of the Feynman diagrams we have used two completely independent approaches, one based on QGRAF [239] and a second one on FEYNARTS [161]. Two independent in-house sets of routines, written in FORM [240] and MATHEMATICA [241] have been used to perform the Lorentz and Dirac algebra. After these first manipulations, the intermediate expressions contain a large number of scalar Feynman integrals, which can be reduced to a much smaller set of MIs by means of IBP and Lorentz Invariant identities [242]. We have executed the reduction algorithms with two independent in-house programs based on the public codes KIRA [112], LITERED [115, 243] and REDUZE2 [117, 244]. We have cross-checked the corresponding expressions at one- and two-loop level, obtained in the two independent approaches, finding perfect agreement.

The evaluation of the quark wave function corrections (Figure 5.2-b) deserves a comment because it proceeds in two steps. First, we generate all the Feynman diagrams contributing to the fermion self-energy and reduce their Feynman integrals to MIs. Second, since we have to compute the derivative of the self-energy with respect to the external invariant, we express the derivative of the MIs as a combination of MIs with a standard reduction procedure. We eventually evaluate this combination in the on-shell limit.

To simplify the calculation, we consider an approximation of the scattering amplitude in the small lepton mass limit, i.e. when m_ℓ is negligible compared to the masses of the gauge bosons and to the energy scales of the process. We thus consider the ratio $m_\ell/\sqrt{s}$ and keep only terms enhanced by a factor $\log(m_\ell^2/s)$, while we discard all those contributions vanishing in the $m_\ell \rightarrow 0$ limit. This approximation is used in the reduction to MIs procedure, in order to distinguish the integrals where the dependence on m_ℓ has to be considered from those where the massless lepton limit can be immediately taken, with a remarkable simplification in the latter case. After such procedure, the only contributions with an explicit leftover dependence on the lepton mass are the ones stemming from factorisable contributions with initial-state QCD and one-loop final-state vertex corrections. Among the box diagrams, those with one or two photons exchanges individually develop logarithms of the lepton mass, but the latter eventually cancel [245] in the physical amplitude.

The identification of all the MIs which contribute to the final result is given using the notion of integral family, which we present in the next Section.

5.1.4 Integral families and the representation of the result in terms of MIs

By associating a scalar integral to a specific integral family, we can express all the scalar products in terms of the denominators of that particular family. We thus obtain, after all the algebraic simplifications, integrals with the following form:

$$I = \int \frac{d^d k_1}{(2\pi)^d} \frac{d^d k_2}{(2\pi)^d} \frac{1}{[k_1^2]^{\alpha_1} [k_2^2]^{\alpha_2} [(k_1 - k_2)^2 - \mu_Z^2]^{\alpha_3} [(k_2 - p_1)^2]^{\alpha_4} \dots}$$

$$\equiv \int \frac{d^d k_1}{(2\pi)^d} \frac{d^d k_2}{(2\pi)^d} \frac{1}{\mathcal{D}_1^{\alpha_1} \mathcal{D}_2^{\alpha_2} [\mathcal{D}_{12} - \mu_Z^2]^{\alpha_3} \mathcal{D}_{2;1}^{\alpha_4} \dots}, \quad (5.29)$$

where the exponents α_k can be either positive (0, 1 or 2) or negative. We introduce the symbols $\mathcal{D}$ s

$$\mathcal{D}_i = k_i^2, \mathcal{D}_{ij} = (k_i - k_j)^2, \mathcal{D}_{i;j} = (k_i - p_j)^2,$$

$$\mathcal{D}_{i;jl} = (k_i - p_j - p_l)^2, \mathcal{D}_{ij;l} = (k_i - k_j - p_l)^2 \quad (5.30)$$

and we use them to precisely define the various integral families. Those relevant for the full calculation are listed here below.

In the cases where the lepton masses are required, because of the presence of final-state collinear singularities, we find the following integral families which accommodate all the appearing scalar integrals:

$$\begin{aligned} B_0 &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_1 &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{12;2}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{20} &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3} - m_\ell^2, \mathcal{D}_{2;3} - m_\ell^2\} \\ B_{20p} &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12}, \mathcal{D}_{1;2}, \mathcal{D}_{2;2}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3} - m_\ell^2, \mathcal{D}_{2;3} - m_\ell^2\} \\ B_{27} &: \{\mathcal{D}_1, \mathcal{D}_2 - m_\ell^2, \mathcal{D}_{12} - m_\ell^2, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12} - m_\ell^2, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\}. \end{aligned} \quad (5.31)$$

Additional families, where the lepton masses can be immediately discarded, are:

$$\begin{aligned} B_{11} &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12} - \mu_V^2, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{12} &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12}, \mathcal{D}_{1;1} - \mu_V^2, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{13} &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12} - \mu_V^2, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{12;2}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{13p} &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12} - \mu_V^2, \mathcal{D}_{1;2}, \mathcal{D}_{2;2}, \mathcal{D}_{1;12}, \mathcal{D}_{12;1}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{14} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_V^2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{14p} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_V^2, \mathcal{D}_{12}, \mathcal{D}_{1;2}, \mathcal{D}_{2;2}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{16} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_V^2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12} - \mu_V^2, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{16p} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_V^2, \mathcal{D}_{12}, \mathcal{D}_{1;2}, \mathcal{D}_{2;2}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12} - \mu_V^2, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ B_{18} &: \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3} - \mu_V^2, \mathcal{D}_{2;3} - \mu_V^2\}, \end{aligned} \quad (5.32)$$

where V can be either Z or W . The massless two-loop form factor integrals were presented in [135]. The MIs with the massive lepton can be straightforwardly obtained from the planar MIs [136, 137] of NNLO QCD corrections to top quark pair production by replacing the top quark mass with the lepton mass. The form factor type two-loop MIs with one or two massive lines were obtained in [133, 134]. The two-loop

box MIs with one and two massive lines have been computed in [138–140]. The reduction to the MIs allows us to obtain an intermediate representation of the bare results, written as a linear combination of MIs with the corresponding rational coefficients

$$\langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(1,1)} \rangle = \sum_{k=1}^{204} c_k(s, t, \{m_i\}, \epsilon) I_k(s, t, \{m_i\}, \epsilon) \quad (5.33)$$

with both the coefficients c_k and the MIs I_k being dependent on the dimensional regularisation parameter ϵ , the kinematical variables s, t and the masses $\{m_i\}$ of the gauge bosons and fermions. The integrals I_k are all of the form $B_i[\vec{\alpha}_j]$, where i labels the integral family and $\vec{\alpha}_j$ is the set of all the exponents.

5.1.5 Infrared singularities and universal pole structure

The renormalised matrix elements are UV-finite, but still contain IR divergences originating from soft and/or collinear massless partons. These singularities are guaranteed to cancel when combined with the NNLO QCD-EW real and real-virtual corrections and subtracted of the initial-state collinear poles, to obtain an IR-safe observable. Nevertheless, their presence at intermediate stages of an inclusive computation, as it is the case with the results provided in this paper, requires the development of methods to consistently handle them; they must be systematically subtracted, leaving a finite remnant.

While in NLO computations the Catani-Seymour dipole subtraction [155, 156, 246] and FKS subtraction [154] are the two most widely used methods, at NNLO several techniques have been proposed (see e.g. [247] and references therein), each one with a different subtraction procedure for the various contributions to physical observables.

The subtraction of the IR poles from the two-loop matrix elements is in general achieved via a process-independent subtraction operator $\mathcal{I}$, that can be constructed by using the universality of the IR singularity structure of the scattering amplitudes, known in the case of a massless gauge theory up to two-loop level [238, 248–250]. An explicit study was performed in [251, 252] to obtain the IR structure of the two-loop amplitudes for mixed QCD $\otimes$ QED corrections to NC Drell-Yan production considering massless leptons. The case of theories with massive particles has been studied in [253–257]. In particular, the IR structure for one-loop QCD corrections to top quark pair production has been examined in detail in ref. [258], while at two-loop level in refs. [259–261]. This can be appropriately Abelianised [262] to obtain the IR structure in the present case by replacing the top quarks with massive leptons.

Different subtraction operator within different subtraction schemes can in principle differ from each other: anyhow the universality of the IR divergences implies that the subtracted virtual contribution at two loops may differ, at most, by a finite constant. For the sake of definiteness, we illustrate here the subtraction operator used in our calculation, which has been performed in the framework of the q_T -subtraction formalism [157], introduced in Section 4.3.1. However, we stress that our results can be easily converted and combined to the description of real radiation in any other IR subtraction scheme, by adding the appropriate finite remainder.

The IR subtraction functions ($\mathcal{I}$) at one-loop are given by

$$\frac{1}{S_\epsilon} \mathcal{I}^{(1,0)} = \left(\frac{s}{\mu^2} \right)^{-\epsilon} C_F \left(-\frac{2}{\epsilon^2} - \frac{1}{\epsilon} (3 + 2i\pi) + \zeta_2 \right), \quad (5.34)$$

$$\frac{1}{S_\epsilon} \mathcal{I}^{(0,1)} = \left(\frac{s}{\mu^2} \right)^{-\epsilon} \left[Q_u^2 \left(-\frac{2}{\epsilon^2} - \frac{1}{\epsilon} (3 + 2i\pi) + \zeta_2 \right) + \frac{4}{\epsilon} \Gamma_l^{(0,1)} \right], \quad (5.35)$$

where

$$\Gamma_l^{(0,1)} = Q_u Q_l \log \left(\frac{2p_1 \cdot p_3}{2p_2 \cdot p_3} \right) + \frac{Q_l^2}{2} \left(-1 - \frac{1+x_l^2}{1-x_l^2} \log(x_l) \right), \quad (5.36)$$

and $S_\epsilon = (4\pi e^{-\gamma_E})^\epsilon$. Q_l and Q_u are the charges of the lepton and of the initial-state quark, and the Casimir of the fundamental representation of $SU(N)$, C_F , is given by $C_F = \frac{N^2-1}{2N}$. The variable x_l is defined by

$$\frac{(1-x_l)^2}{x_l} = -\frac{s}{m_\ell^2}. \quad (5.37)$$

Using the one-loop subtraction functions, we obtain the finite contributions to the one-loop QCD and EW amplitudes, respectively, as follows:

$$|\mathcal{M}^{(1,0),fin}\rangle = |\mathcal{M}^{(1,0)}\rangle - \mathcal{I}^{(1,0)} |\mathcal{M}^{(0,0)}\rangle, \quad (5.38)$$

$$|\mathcal{M}^{(0,1),fin}\rangle = |\mathcal{M}^{(0,1)}\rangle - \mathcal{I}^{(0,1)} |\mathcal{M}^{(0,0)}\rangle. \quad (5.39)$$

The mixed two-loop subtraction operator⁴ is given by

$$\begin{aligned} \frac{1}{S_\epsilon^2} \mathcal{I}^{(1,1)} = & \left(\frac{s}{\mu^2} \right)^{-2\epsilon} C_F \left[Q_u^2 \left(\frac{4}{\epsilon^4} + \frac{1}{\epsilon^3} (12 + 8i\pi) + \frac{1}{\epsilon^2} (9 - 28\zeta_2 + 12i\pi) \right. \right. \\ & + \frac{1}{\epsilon} \left(-\frac{3}{2} + 6\zeta_2 - 24\zeta_3 - 4i\pi\zeta_2 \right) \Big) \\ & \left. + \left(-\frac{2}{\epsilon^2} - \frac{1}{\epsilon} (3 + 2i\pi) + \zeta_2 \right) \frac{4}{\epsilon} \Gamma_l^{(0,1)} \right]. \end{aligned} \quad (5.40)$$

In order to describe the contributions entering in $\mathcal{I}^{(1,1)}$, let us discuss first the IR singularities of the one-loop and two-loop boxes. The fact that the $q\bar{q}$ initial state has zero electric charge implies that there are no QED collinear singularities stemming from the exchange of a photon between an initial quark and a final lepton lines. Such terms appear in the individual diagrams, but cancel in the physical amplitude. This conclusion can also be seen as a confirmation of the universality of initial state QED collinear singularities. This property is observed at one loop and confirmed also when going to $\mathcal{O}(\alpha\alpha_s)$ by dressing the initial state quark line with all the possible gluon exchanges. In the sum of all the two-loop box diagrams we can thus expect that the photon exchange develops at most a soft divergence. The presence of a colorful initial state quark line but of a colorless final state lepton line leads to a natural separation between the gluon IR-divergent exchanges in the initial state and the QED interaction in the rest of the diagram. For this reason, all the IR singularities, soft and collinear, stemming from the two-loop box diagrams, can be subtracted by the product of $\mathcal{I}^{(1,0)}$ and $\mathcal{I}^{(0,1)}$, times the Born amplitude. The divergences in the product of initial and final factorizable contributions are obviously canceled by the same product of one-loop subtraction operators. The initial state vertex corrections have a richer structure, including planar and non-planar topologies, both yielding IR divergent contributions. The former can be subtracted again by the product of $\mathcal{I}^{(1,0)}$ and $\mathcal{I}^{(0,1)}$, while the latter require a genuinely two-loop subtraction operator. From the above comments we conclude that the operator $\mathcal{I}^{(1,1)}$ is applied to the Born amplitude and receives contributions from the product of $\mathcal{I}^{(1,0)}$ and $\mathcal{I}^{(0,1)}$ and from the two-loop non-factorizable contributions appearing in the vertex corrections.

⁴An analogous expression, in the massless lepton case, has been discussed in refs. [251, 252].

We note that the application of $\mathcal{I}^{(1,1)}$ can not make the complete amplitude IR finite. There are in fact additional configurations where e.g. the gluon yields a divergence and the EW interaction is finite. This is for instance the case in the two-loop boxes with the exchange of two W or Z bosons. They can be subtracted by applying the one-loop QCD subtraction operator $\mathcal{I}^{(1,0)}$ to the one-loop EW amplitude subtracted of its one-loop QED divergences, as defined in Eq. 5.39. An analogous subtraction can remove the photon divergences which multiply a finite QCD remnant, as defined in Eq. 5.38.

Using Eqs. 5.38-5.40, we obtain the finite and subtracted two-loop amplitude as follows

$$|\mathcal{M}^{(1,1),fin}\rangle = |\mathcal{M}^{(1,1)}\rangle - \mathcal{I}^{(1,1)}|\mathcal{M}^{(0,0)}\rangle - \tilde{\mathcal{I}}^{(0,1)}|\mathcal{M}^{(1,0),fin}\rangle - \tilde{\mathcal{I}}^{(1,0)}|\mathcal{M}^{(0,1),fin}\rangle, \quad (5.41)$$

with $\tilde{\mathcal{I}}^{(i,j)}$ defined by dropping the term proportional to ζ_2 in $\mathcal{I}^{(i,j)}$. This choice is conventional and defines the finite part of our subtraction term. The approximation of the amplitude in the small lepton mass limit retains all the terms enhanced by $\log(m_\ell)$, divergent in the $m_\ell \rightarrow 0$ limit. The structure of these corrections reflects the universality property of the final-state collinear divergences, and is given, normalised to the Born squared matrix element, by

$$\lim_{m_\ell \rightarrow 0} \frac{\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1),fin} \rangle}{\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,0)} \rangle} = K + C_F Q_l^2 (-8 + 7\zeta_2 - 3i\pi) \left[-\log\left(\frac{m_\ell^2}{s}\right) + \log^2\left(\frac{m_\ell^2}{s}\right) \right], \quad (5.42)$$

where K represents all the other terms in the interference, constant in the $m_\ell \rightarrow 0$ limit.

5.1.6 The UV-renormalised IR-subtracted hard function

Using Eq. 5.41, we obtain the finite IR-subtracted matrix element $\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1),fin} \rangle$ from the UV renormalised matrix element $\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1)} \rangle$

$$\begin{aligned} \langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1),fin} \rangle &= \langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1)} \rangle - \mathcal{I}^{(1,1)} \langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,0)} \rangle \\ &\quad - \tilde{\mathcal{I}}^{(0,1)} \langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,0),fin} \rangle - \tilde{\mathcal{I}}^{(1,0)} \langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,1),fin} \rangle. \end{aligned} \quad (5.43)$$

This quantity, renormalised with the Born squared matrix element, represents the hard function $H^{(1,1)}$, defined as

$$H^{(1,1)} = \frac{1}{16} \left[2 \operatorname{Re} \left(\frac{\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(1,1),fin} \rangle}{\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,0)} \rangle} \right) \right]. \quad (5.44)$$

The contributions to the two-loop vertex and box diagrams, introduced in Eq. 5.33, come from 204 MIs. We compute them using the method of differential equations [23, 242, 263–267]. For a large subset, we express the solution in terms of generalised harmonic polylogarithms (GHPLs) [268–270]. In our results we also find the following 5 MIs:

$$\begin{aligned} &\{0, 1, 1, 1, 0, 1, 1, 0, 1\}, \{1, 1, 1, 1, 0, 1, 1, 0, 1\}, \{1, 1, 1, 1, -1, 1, 1, 0, 1\}, \\ &\{1, 1, 1, 1, 0, 1, 1, -1, 1\}, \{1, 1, 1, 1, -1, 1, 1, -1, 1\}, \end{aligned} \quad (5.45)$$

in the two integral families B_{16} and B_{16p} which are not expressed, according to the discussion of ref. [138], in terms of GHPLs, but using Chen-Goncharov integrals [141]. The possibility to solve the MIs with two internal massive lines and obtain a representation in terms of Nielsen polylogarithmic functions has been discussed in ref. [139]. In our case, the evaluation of the Chen iterated integrals [141] in a non-physical region, where the boundary conditions are fixed, is possible, but the analytic continua-

tion to cover the full physical region constitutes a formidable task. For this reason, we adopt a different, semi-analytical, approach to solve and evaluate these MIs, and we describe it in Section 5.1.7.

The final results, expanded in powers of ϵ , are written in terms of GHPLs and of the symbols associated to the 5 MIs whose solution is given in terms of Chen integrals. In the next Sections, we explicitly discuss the numerical evaluation of the latter. We note that the individual contributions from these 5 MIs, to the single pole in ϵ of the full unrenormalised matrix element, contain the Chen iterated integrals, but the summed result is independent of those.

5.1.7 Semi-analytical solution of MIs via series expansion

In ref. [138], 5 MIs with two internal massive lines have been solved in terms of Chen iterated integrals, but their analytical continuation to the physical region is extremely difficult. We thus turn to a different approach and we obtain these MIs in semi-analytical form, by solving the relevant system of differential equations via series expansion as initially proposed in ref. [132]. This approach is implemented in the MATHEMATICA code DIFFEXP [106]. However, since the public version of DIFFEXP does not support, for the time being, the possibility of complex masses in the internal lines of the loop integrals, we relied on our package SEASYDE [2], that consistently takes into account the presence of complex weights and the analytic continuation to the complex plane. By approaching the problem in such a semi-analytical way, we have to consider the complete system (36×36) of differential equations [138], which includes the 5 MIs of Eq. 5.45. The other 31 integrals are already known in terms of GHPLs and we use their numerical evaluation obtained with GiNAC [271] as a cross check of the correct behavior of SEASYDE finding perfect agreement⁵. The values for the 5 MIs in Eq. 5.45 are instead new predictions. In the Euclidean region, we validated our results for these five MIs against the values computed with PYSECDEC [129]. We found perfect agreement up to the first five significant digits, which corresponds to the numerical precision limit achieved by PYSECDEC. This evaluation required approximately two hours of computation time on a standard laptop. In the physical region, however, a similar comparison could not be performed due to convergence issues encountered with PYSECDEC, leading us to adopt an alternative validation strategy. We have solved the same system of differential equations with a real value for the gauge-boson masses, by using both DIFFEXP and SEASYDE, finding perfect agreement⁶, for different values of the centre-of-mass energy, below and above all the physical thresholds. Finally, we cross-checked against AMFLOW the values of the five MIs with complex internal masses, in the physical region. We found exact agreement up to 40 significant digits. This numerical evaluation required 50 minutes of computation time on a 32-core machine.

In Figure 5.5 we present the real and imaginary parts of the values of the ϵ^0 coefficient of the MIs numbered from 32 to 36 in ref. [138], in a phase-space region relevant for LHC phenomenology.

This approach faces all the problems of the analytical continuation immediately from the analysis of the structure of the differential equations. If it is possible to expand the solution in a region around the boundary condition points, then the problem becomes of algebraic nature, i.e. we have to solve for all the coefficients of the various powers of the expansion variable. The system of differential equations of ref. [138] can be cast in dlog form and thus contains the complete information about the singularity structure of the problem, in particular the position of the branch points relevant in the discussion of the polylogarithmic functions present in the solution. This information can be exploited to control the

⁵Both codes work with arbitrary precision and any level of agreement can be reached simply adding more terms in the respective expansions. In particular we have verified explicitly agreement of the first 40 significant digits

⁶Also in this case the two codes can work with arbitrary precision and we have explicitly checked the agreement of the first 40 digits.




Figure 5.5: Real (left column) and imaginary (right column) parts of the ϵ^0 coefficient of the two-mass MIs labeled 32 to 36 in ref. [138].

analytic continuation when we extend the solution to an adjacent region, either non-physical or physi-




Figure 5.6: In the plot we show the finite hard function $H^{(1,1)}$, in the resonance region (left) and in the high invariant mass region (right), as a function of $\sqrt{s}$ and $\cos \theta$.

cal. We remark that complete control on the analytic continuation has been verified when varying one kinematical variable at the time.

5.1.8 Numerical results

In this Section, we present the numerical evaluation of our result, the finite IR-subtracted UV-renormalised hard function $H^{(1,1)}$. For every practical application, the evaluation in one phase-space point of the correction factor stemming from the two-loop $\mathcal{O}(\alpha\alpha_s)$ virtual corrections requires a cumbersome procedure. We consider it more convenient to prepare a numerical grid for $H^{(1,1)}$, covering all the phase-space values relevant for NC DY in a given fiducial volume and parameterised in terms of the partonic centre-of-mass energy $\sqrt{s}$ and scattering angle $\cos \theta$. We notice that the virtual correction, after the subtraction of all the IR singularities, is a smooth, slowly varying function of $\sqrt{s}$ and $\cos \theta$. We have prepared a grid with a sampling based on the known behaviour of the NLO-EW distribution, with special care for the Z resonance region, where a finer binning is necessary. For illustration purposes we consider here a grid with (130x25) points in $(\sqrt{s}, \cos \theta)$, covering the intervals $s \in [40, 13000]$ GeV and $\cos \theta \in [-1, 1]$. The extraction of any intermediate value of $H^{(1,1)}$ can be obtained via interpolation of the grid points with excellent accuracy thanks to the smoothness of the function. With our implementation, we reproduce the exact correction with an accuracy at the 10^{-3} level, in the whole phase space. This accuracy is sufficient for phenomenological studies, but can be increased whenever needed.

The preparation, on a single core, of a grid of values for all the 36 two-masses MIs requires a fixed step, executed only once, to move from the boundary conditions fixed in the non-physical Euclidean region to the physical region, which can take ~ 20 minutes. Moving in the physical region is faster, with an average time of ~ 5 minutes per point. Once this first grid is ready, the evaluation of the corresponding grid for the $H^{(1,1)}$ function requires ~ 3 minutes per point. Once the $H^{(1,1)}$ grid is ready, its interpolation in any simulation code requires a negligible time.

We use the following parameters:

m_Z	91.1535 GeV	Γ_Z	2.4943 GeV
m_W	80.358 GeV	Γ_W	2.084 GeV
m_H	125.09 GeV	m_t	173.07 GeV

The values of the gauge boson masses and decay widths are compatible with those reported in the PDG [272], extracted in a running-width fitting scheme. The choice of the complex mass scheme to define the masses of the unstable W and Z bosons implies that all the functions which depend on these parameters become functions of complex variables. We have exploited the rescaling properties of the GHPLs to

set their argument equal to one, adjusting all the other weights, which become in general complex⁷. For the numerical evaluation of all the GHPLs appearing in the final expressions, we use GINAC and HANDYG [142] in two independent C++ and MATHEMATICA codes to evaluate each phase-space point. We also use HARMONICSUMS [143, 144], POLYLOGTOOLS [145] and LOOPTOOLS [146] for several cross-checks of the one- and two-loop integrals available in closed analytic form, with real and complex masses. For the numerical evaluation of the 5 MIs of Eq. 5.45, we use the semi-analytical approach described in Section 5.1.7.

We present in Figure 5.6 the numerical grid expressing the UV-renormalised IR-subtracted two-loop $\mathcal{O}(\alpha\alpha_s)$ virtual correction due to the two-loop vertex and box diagrams (Figure 5.2 a,b,f and Figure 5.3 a,b,f). The plots show the IR subtracted correction factor $H^{(1,1)}$ defined in Eq. 5.44, as a function of $\sqrt{s}$ and $\cos\theta$. We consider the range $-1 \leq \cos\theta \leq 1$ and, for the partonic center-of-mass energy, two different intervals, namely $50 \leq \sqrt{s} \leq 250$ GeV and $50 \leq \sqrt{s} \leq 5000$ GeV.

With the zoom in the $50 \leq \sqrt{s} \leq 250$ GeV interval, we can appreciate the non-trivial structure of the corrections starting from the Z resonance and up to 110 GeV. The description of the WW and ZZ thresholds is smooth, thanks to the implementation of the complex-mass scheme. Above 200 GeV the shape of the correction factor is negative, smooth, and slowly decreasing; we do not observe the appearance of additional structures. The non-trivial $\cos\theta$ dependence is due to parity violating effects, in largest size stemming from the interference between the two-loop AA subset of diagrams and the tree-level amplitude with a Z -boson exchange.

5.2 Two-loop mixed corrections to CC-DY

The process under study is the production of a lepton-neutrino pair in quark-antiquark annihilation

$$u(p_1) + \bar{d}(p_2) \rightarrow \nu_\ell(p_3) + \ell^+(p_4). \quad (5.46)$$

We consider the two-loop mixed QCD-EW corrections to this reaction⁸, with explicit results for massless initial-state quarks. The Mandelstam variables are defined as:

$$s = (p_1 + p_2)^2, \quad t = (p_1 - p_3)^2, \quad u = (p_2 - p_3)^2 \quad \text{with } s + t + u = m_\ell^2, \quad (5.47)$$

while the on-shell conditions of the external particles are:

$$p_1^2 = p_2^2 = p_3^2 = 0; \quad p_4^2 = m_\ell^2; \quad (5.48)$$

5.2.1 Generation of the full amplitude and evaluation of the interference terms

Different kinds of Feynman diagrams contribute, at $\mathcal{O}(\alpha\alpha_s)$, to the scattering amplitude, and we compute them using the background field gauge (BFG) [48]. We present⁹ in Figures 5.7 and 5.8 a few representative examples. The initial state two-loop vertex corrections (Figure 5.7-(a)), are combined with the two-loop external quark wave function corrections (Figure 5.7-(b)) and with the one-loop external quark wave function correction with initial-state QCD vertex correction (Figure 5.8-(f)). The combination yields an UV-finite, but still IR-divergent, result. We treat as a separate subset the sum of two-loop

⁷Cfr. also ref. [273]

⁸The production of the $\ell^- \bar{\nu}_\ell$ final state can be obtained with a fully analogous calculation, or, assuming CP invariance, from the amplitude for the positively charged final state, with appropriate transformations of the external momenta.

⁹The straight line with an arrow represents a fermion line (blue: quark, red: leptons), the wavy lines represent EW gauge bosons (green: W , red: Z , black: photon) and the spiral coils represent gluons.




Figure 5.7: Sample Feynman diagrams of two-loop corrections and associated two-loop counterterms. The wavy green line represent a W -boson while the black one a photon.




Figure 5.8: Sample Feynman diagrams of factorisable corrections, including one-loop counterterm corrections. The wavy green line represent a W -boson while the red one a Z -boson.

W self-energy, together with the corresponding two-loop mass counterterms and with the insertion of charge renormalisation constants in the initial and final state vertices. Representative diagrams are shown in Figures 5.7-(c),(d),(e). The combination yields an UV- and IR-finite contribution. An example of a two-loop box with the exchange of one W and one neutral EW boson is given in Figure 5.7-(f). We note the absence of contributions from closed fermionic triangles due to colour conservation. Specifically, at this perturbative order, such a loop would necessarily connect to the remainder of the diagram via exactly one gluon and two electroweak gauge bosons. Because the electroweak bosons are color-neutral, the associated color factor is proportional to the trace of a single $SU(3)_c$ generator, which vanishes identically: $\text{Tr}(t^a) = 0$.

The factorisable contributions are schematically represented in Figures 5.8-(a),(f). They all include, at $\mathcal{O}(\alpha\alpha_s)$, an initial-state QCD vertex correction. The second factor can be, alternatively: the final-state EW vertex corrections, the external lepton wave function correction, the one-loop W self-energy corrections, the one-loop W mass renormalisation counterterms, the one-loop external quark wave function correction and the one-loop charge renormalisation contributions. The properties of the BFG allow to identify UV-finite Feynman diagrams combinations.

We follow the general procedure introduced in Section 4.4.1 to compute the bare matrix elements up to two-loop order.

We retain the exact dependence on the lepton mass m_ℓ , both at the one- and two-loop level, only in the diagrams where an internal photon couples to the final-state lepton, featuring the complete set of contributions enhanced by $\log(m_\ell^2)$. The lepton mass regulates, in this case, the final-state collinear divergences. These diagrams also yield a gauge-invariant subset of terms proportional to m_ℓ^2/m_W^2 , which are however phenomenologically negligible. In all the other cases, we do not have final-state IR divergences, and therefore we can keep the lepton massless. The use of this approximation reduces by one the number of independent mass scales in the computation and thus its complexity, allowing for a faster evaluation of the amplitude.

5.2.2 Integral families and the representation of the result in terms of MIs

The scattering amplitude, after the computation of the interference terms, appears as a sum of scalar Feynman integrals.

In order to classify all the integrals entering in our computation, we need 11 different integral families. Among them, 7 can be defined exactly as the ones introduced in Sec. 5.1 for the case of the NC DY process:

$$B_0, B_1, B_{11}, B_{12}, B_{13}, B_{14}, B_{18}. \quad (5.49)$$

The remaining 4 integral families only appear for the case of CC DY:

$$\begin{aligned} \tilde{B}_{14} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_V^2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3} - m_\ell^2\} \\ \tilde{B}_{14p} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_V^2, \mathcal{D}_{12}, \mathcal{D}_{1;2}, \mathcal{D}_{2;2}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12}, \mathcal{D}_{1;3}, \mathcal{D}_{2;3} - m_\ell^2\} \\ \tilde{B}_{16} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_{V_1}^2, \mathcal{D}_{12}, \mathcal{D}_{1;1}, \mathcal{D}_{2;1}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12} - \mu_{V_2}^2, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\} \\ \tilde{B}_{16p} &: \{\mathcal{D}_1, \mathcal{D}_2 - \mu_{V_1}^2, \mathcal{D}_{12}, \mathcal{D}_{1;2}, \mathcal{D}_{2;2}, \mathcal{D}_{1;12}, \mathcal{D}_{2;12} - \mu_{V_2}^2, \mathcal{D}_{1;3}, \mathcal{D}_{2;3}\}, \end{aligned} \quad (5.50)$$

where V_1 and V_2 are two different vector bosons and the $\mathcal{D}$ s are the ones introduced in Eq. 5.30.

The integral families listed in Eq. 5.50, despite being specific to the charged current case, can also be seen as an extension of integral families already appearing in the neutral current process, where the value of the masses in the propagators has been modified. In particular, $\tilde{B}_{16}$ describes box integrals with a W - Z exchange and is the natural extension of B_{16} , already introduced in [1] to describe the Z - Z and W - W box integrals. The latter can in fact be obtained from $\tilde{B}_{16}$ in the equal mass limit $\mu_{V_1} \rightarrow \mu_V$, $\mu_{V_2} \rightarrow \mu_V$. Similarly, by considering a massless lepton, $m_\ell = 0$, in the integral family $\tilde{B}_{14}$, we retrieve the integral family B_{14} listed in Eq. (5.49).

In the NC DY case, we introduced the integral family B_{14} instead of the more general $\tilde{B}_{14}$ because the final state collinear singularities, like e.g. those stemming from the γ - Z boxes, were cancelling between different contributions [245]. We were thus able to take the massless limit for this subset of diagrams. The integral family B_{14} yields collinear poles in dimensional regularization, while such singularities appear in $\tilde{B}_{14}$ as lepton-mass logarithms.

After performing the reduction to MIs, the bare result can be cast in the following, compact form:

$$\langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(1,1)} \rangle = \sum_{k=1}^{274} c_k(s, t, \{m_i\}, \epsilon) I_k(s, t, \{m_i\}, \epsilon) \quad (5.51)$$

where c_k are rational coefficients, functions of the kinematical invariants (s, t) , the set of real and complex masses $\{m_i\}$ and of the regularisation parameter ϵ , while I_k are the MIs¹⁰. In addition to the MIs

¹⁰This intermediate result is provided as an ancillary file attached to this paper.




Figure 5.9: Master Integrals with two different internal masses, belonging to topology $\tilde{B}_{16}$. Thin plain lines represent massless particles. Thick and double lines represent massive particles. External dashed lines indicate the Mandelstam invariant s .




Figure 5.10: Master Integrals with two different masses, belonging to topology $\tilde{B}_{14}$. Thin plain lines represent massless particles. Thick and double lines represent massive particles. External dashed lines indicate the Mandelstam invariant s .

employed for the NC DY, in the case of the CC DY we have to evaluate two more sets. In Figure 5.9 we show the additional MIs belonging to topology $\tilde{B}_{16}$, that, therefore, have two different masses in the propagators. In Figure 5.10, instead, we show the MIs of topology $\tilde{B}_{14}$, that have two different masses in the propagators and a massive external leg.

5.2.3 Numerical evaluation of MIs via series expansions

The final step of the calculation consists in the numerical evaluation of the MIs. The additional complication of the charged-current calculation, with respect to the neutral-current one presented in [1], is given by the presence of the $\tilde{B}_{14}$ and $\tilde{B}_{16}$ topologies, that are two-loop boxes with two internal (and one external in the case of $\tilde{B}_{14}$), and different, massive lines. In the case of $\tilde{B}_{16}$, the analytical result for the equal masses case has been presented in [138] in terms of Chen-iterated integrals and an analytical expression in terms of GPLs has been found in [139]. However, for the different masses case, no analytical expressions are available in the literature. We have opted to compute the integrals using the series expansion approach.

We have generated the differential equations using in-house MATHEMATICA routines, which rely on the external package LITERED for computing the derivatives with respect to kinematic invariants s and t and on the IBPs relations provided by KIRA. The boundary conditions for all the MIs have been computed in a point in the physical region by using AMFLOW [130], interfaced with KIRA. Their numerical evaluation, up to 50 digits, requires ~ 4.5 h on a laptop using 8 threads. Finally, we have solved the system of differential equations in the Mandelstam variables s and t using the MATHEMATICA package SEASYDE [2], that allows us to consider complex-valued internal masses.

Since when combining the numerical values of the MIs with their rational coefficients big numerical cancellations can occur, it is extremely important to keep under control the numerical precision of the evaluation of the MIs. Within SEASYDE, we have decided to keep 70 terms in the expansion in the kinematic invariants. The estimate for the relative error provided by SEASYDE is, at worst, 10^{-14} in every point of the phase-space¹¹. In several points of the phase-space we have performed a cross-check against AMFLOW by computing the MIs with 50 digits precision, finding agreement for the first 14 digits, in accordance to the uncertainty estimated by SEASYDE. Cross-checks with analytical expressions, when available, have been performed as-well, finding agreement.

Using SEASYDE we were able to generate a numerical grid in $(\sqrt{s}, \cos \theta)$, consisting in 3250 points, as described in detail in Section 5.1.8. The numerical evaluation of the complete grid for all the topologies requires, in total, ~ 3 weeks on a cluster with 26 cores. The most complicated case is the two-loop box with two massive internal lines, for which we have to solve a system with 56 equations. The calculation of its grid requires ~ 10 days by itself.

Mass evolution

Finally, we have done a self-consistency check by exploiting the flexibility given by the method of differential equations and the series expansion approach. We observe, indeed, that in the limit of the two masses being equal, the $\tilde{B}_{16}$ topology reduces to the B_{16} from the neutral-current case. Hence, another way for computing a numerical grid for $\tilde{B}_{16}$ is, firstly, to create a numerical grid in (s, t) for B_{16} and, secondly, to write down the differential equations w.r.t. one of the two masses to evolve B_{16} to $\tilde{B}_{16}$. We

¹¹To estimate the numerical precision of the series expansion in SEASYDE, we evaluate the n -term series at half its radius of convergence and compare it against the same expansion truncated at $n - 3$ terms. The relative difference between these two evaluations is taken as the error estimate.




Figure 5.11: The two possibilities for creating a numerical grid for the $\tilde{B}_{16}$ topology. The first one, represented by the blue line, is to get the BCs for $\tilde{B}_{16}$ and then use the diff. eqs. w.r.t. s and t for obtaining the final grid. The second one, indicated by the orange lines, is to obtain the BCs for B_{16} , use the diff. eqs. for creating a grid in s and t for B_{16} and, finally, use it as new BCs for evolving B_{16} to $\tilde{B}_{16}$.

have explicitly verified, for different points in the phase-space, that the result obtained by evolving the equal masses box in s and t and then in one of the two masses, is in agreement with the one obtained by evolving directly the two-masses box in the Mandelstam variables s and t . The two possibilities are schematically depicted in Figure 5.11.

5.2.4 Additional comments on the numerical evaluation of MIs

In this work, our main objective is to pursue the automation of all the steps in the calculation of virtual corrections, in particular, the numerical evaluation of the MIs. Within our framework, all the steps are handled automatically (computing the differential equations, obtaining the boundary conditions and solving the system) thus opening the door to more complicated problems with the presence of a larger number of topologies.

However, this automation comes at a price. In [1], indeed, the computation of a numerical grid for the B_{16} topology takes ~ 12 hours. Even considering a factor $1.5/2$ related to the bigger size of the system (56 equations for $\tilde{B}_{16}$ against 36 for B_{16}), we are a factor $10/15$ slower with respect to the case of NC DY. This big difference in time is mainly due to the state of the system of differential equations. In the previous case, in fact, we have put a strong effort in simplifying the system by choosing a suitable change of basis to cast the system in pre-canonical form. This simpler form of the system enormously speeds up the computation. On the contrary, in this work, we have taken the MIs suggested by KIRA and we have blindly written down their differential equations. This means that, other than standard algebraic simplifications, we have not performed any other optimization. Finding a change of basis which simplifies the system of differential equations is indeed important. However, it requires a big effort and a profound knowledge of the problem, while this blind approach relies only on computational power.

A second comment applies to the possibility of using AMFLOW for the evaluation of the MIs in all the points of the phase-space. The evaluation of all the MIs takes $\sim 3\text{h}15\text{m}$ per point, running on 8 cores and asking for a precision of 16 digits. By using the same setup as the one used for SEASYDE, the complete run would approximately take 130 days¹².

¹²The run-times have been estimated by using the latest public versions of each code.

A final comment regards a severe limiting factor when working with numerical grids. Usually, all the input parameters are fixed to their numerical values and they cannot be modified without rerunning the final steps of the calculation. This requires a large amount of time. The series expansion approach, however, can be exploited for providing, with limited additional work, also the dependence on an additional variable e.g. $\delta m_W \equiv m_W - \overline{m}_W$, where $\overline{m}_W$ is a reference value for the mass of the W boson. This can be done following the idea of the mass evolution introduced in Section 5.2.3. In particular, for each topology, we use the grid in (s, t) , that we already have with a reference value $\overline{m}_W$, as a boundary condition, and we solve the differential equations w.r.t. m_W around $\overline{m}_W$. By doing so, it is possible to provide a grid in (s, t) where each point is a series in δm_W . Moreover, since the δm_W we are interested in are of the order of 0.1 GeV at most, only a small number of terms in the series is needed. Therefore, the solution can be computed in a small amount of time. For illustration, the evaluation of the δm_W dependence for the $\tilde{B}_{16}$ topology, for all 3250 points, has taken ~ 1.5 days, keeping 15 terms in the series. The difference between the SEASYDE solution, evaluated in $\delta m_W = 0.1$ GeV, and AMFLOW is on the 13th significant digit. This approach opens the possibility for implementing this kind of corrections in frontier studies like, but not limited to, the W -boson mass determination.

5.3 Two-loop EW corrections to NC-DY: IR structure

Now we also present the evaluation of the two-loop pure QED corrections to NC-DY. This QED calculation serves as a strategic, intermediate step toward the ultimate goal of a full two-loop electroweak evaluation. From a technical point of view, the pure QED case offers a highly instructive framework, since it significantly reduces the number of Feynman diagrams and the complexity of Feynman Integrals due to the absence of massive gauge bosons. Most importantly, the IR pole structure of the pure QED result is the same as the one of the full two-loop EW case. Consequently, successfully isolating and subtracting the IR divergences within the QED amplitude acts as a rigorous proof of concept, demonstrating that the applied subtraction methodology will successfully and extend to the full two-loop EW calculation.

5.3.1 Infrared singularities and universal pole structure

The IR divergent structure of the bare amplitude $|\hat{\mathcal{M}}\rangle$ can be captured in a process-independent way by exploiting the universal structure of IR divergences [274]. This is made possible by the introduction of suitable subtraction operators $\mathcal{I}^{(i,j)}$, analogous to the ones introduced in previous Sections for the mixed corrections. In particular, the operators $\mathcal{I}^{(0,i)}$ in EW/QED, describe the emission of massless photons and the production of additional light fermion pairs. The scattering amplitude diverges when the emitted particles have vanishing energy (soft limit) or form a vanishing angle with their emitter (collinear limit). We regularize the divergences arising from the integration over the loop momenta by working in $d = 4 - 2\epsilon$ space-time dimensions. In the case of emissions off the massive final state, the emitter mass m acts as a natural regulator, the collinear divergence represented by a $1/\epsilon$ pole is absent, but in the amplitude we find logarithms of m . In the case of EW corrections, the conservation of the vector current implies that the impact of genuinely weak corrections exactly vanishes in the divergent emission factors, whose expression receives only QED contributions. The universality of the IR divergent factors implies that their expression depends only on the identity and the kinematics of the external particles of the scattering under study, but not on the details of the hard interaction; in other words, the discussion of the IR structure of the $u\bar{u} \rightarrow e^+e^-$ process is the same, irrespective of the fact that we have a QED or an EW interaction driving the scattering process.

Considering the case of QED corrections up to second order, we have:

$$|\mathcal{M}^{(0,1),fin}\rangle = |\mathcal{M}^{(0,1)}\rangle - \mathcal{I}^{(0,1)}|\mathcal{M}^{(0,0)}\rangle \quad (5.52)$$

$$|\mathcal{M}^{(0,2),fin}\rangle = |\mathcal{M}^{(0,2)}\rangle - \mathcal{I}^{(0,2)}|\mathcal{M}^{(0,0)}\rangle - \mathcal{I}^{(0,1)}|\mathcal{M}^{(0,1),fin}\rangle. \quad (5.53)$$

We restrict the discussion to the pure QED corrections to a scattering process $2 \rightarrow n$, where the final state only contains massive emitters. In this case, the infrared structure of the amplitude is encapsulated in the subtraction operators $\mathcal{I}^{(0,1)}$ and $\mathcal{I}^{(0,2)}$, defined as

$$\begin{aligned} \frac{1}{S_\epsilon} \mathcal{I}^{(0,1)} &= \frac{\Gamma'_0}{4\epsilon^2} + \frac{\Gamma_0}{2\epsilon}, \\ \frac{1}{S_\epsilon^2} \mathcal{I}^{(0,2)} &= \frac{(\Gamma'_0)^2}{32\epsilon^4} + \frac{\Gamma'_0}{8\epsilon^3} \left(\Gamma_0 - \frac{3}{2}\beta_0 \right) + \frac{\Gamma_0}{8\epsilon^2} (\Gamma_0 - 2\beta_0) + \frac{\Gamma'_1}{16\epsilon^2} + \frac{\Gamma_1}{4\epsilon} \\ &\quad + \sum_k^{\text{massive}} b_k n_k q_k^2 \left\{ \frac{1}{4\epsilon^2} \Gamma'_0 \log \left(\frac{\mu_R^2}{m_k^2} \right) + \frac{1}{8\epsilon} \left[\Gamma'_0 \left(\log^2 \left(\frac{\mu_R^2}{m_k^2} \right) \right. \right. \right. \\ &\quad \left. \left. \left. + \frac{\pi^2}{6} \right) + 4\Gamma_0 \log \left(\frac{\mu_R^2}{m_k^2} \right) \right] \right\}, \end{aligned} \quad (5.54)$$

where $S_\epsilon = (4\pi e^{-\gamma_E})^\epsilon$, μ_R is the renormalisation scale and the symbols $\Gamma_i^{(l)}$ will be defined in the following. In Eq. 5.55 the sum runs over the particles considered massive in the computation, with multiplicity n_k (3 for the coloured quarks, 1 for the leptons, 1 for the W boson), electric charge q_k , and mass m_k . The coefficients b_k (with $b_f = -\frac{4}{3}$ in the case of fermions and $b_W = +7$ for the W boson) reflect the expression of the EW β function: $\beta_0^{EW} = 7 - \frac{4}{3} \sum_k n_k q_k^2$. In the case of a pure QED calculation, we must set $b_W = 0$ in Eq. 5.55. The structure of the subtraction operators is governed by the functions Γ and Γ' , which can be obtained from the Abelianization of the QCD soft and cusp anomalous dimensions [255, 275]

$$\Gamma = \sum_i \gamma_i^i + Q_1 Q_2 \gamma^{\text{cusp}}(\alpha) \log \left(\frac{\mu_R^2}{-2p_1 \cdot p_2} \right) \quad (5.56)$$

$$\begin{aligned} &+ \sum_j \gamma_h^j - \sum_{j_1 \neq j_2} \frac{Q_{j_1} Q_{j_2}}{2} \gamma^{\text{cusp}}(\alpha) \beta_{j_1, j_2} \coth \beta_{j_1, j_2} \\ &- \sum_{i, j} Q_i Q_j \gamma^{\text{cusp}}(\alpha) \log \left(\frac{\mu_R m_j}{2p_i \cdot p_j} \right), \end{aligned}$$

$$\Gamma' = - \sum_i Q_i^2 \gamma^{\text{cusp}}(\alpha), \quad (5.57)$$

$$\gamma^{\text{cusp}}(\alpha) = \frac{\alpha}{4\pi} \left[4 - \frac{80}{9} \left(\frac{\alpha}{4\pi} \right) \sum_k^{\text{massless}} n_k q_k^2 + \mathcal{O}(\alpha^3) \right]. \quad (5.58)$$

We refer to massless initial-state particles, with momentum p_i and electric charge Q_i , while with the labels $j = 3, \dots, n$ we refer to the massive final state emitters, with mass m_j , momentum p_j and electric charge Q_j . The first line of Eq. 5.56 thus represents the contribution coming from initial-state radiation,

with γ_l^i defined as

$$\gamma_l^i = \frac{\alpha}{4\pi} Q_i^2 \left\{ -3 + \left(\frac{\alpha}{4\pi} \right) \left[Q_i^2 \left(-\frac{3}{2} + 2\pi^2 - 24\zeta_3 \right) + \left(\frac{130}{27} + \frac{2\pi^2}{3} \right) \sum_k^{\text{massless}} n_k q_k^2 \right] + \mathcal{O}(\alpha^3) \right\}. \quad (5.59)$$

In Eq. 5.60 the sum runs over the particles considered massless in the computation, with multiplicity n_k (3 for the coloured quarks, 1 for the leptons) and electric charge q_k , giving the contribution of the soft and collinear emission of light fermion pairs from the initial-state partons. The second line of Eq. 5.56 instead contains the contribution coming from the final-state particles, with γ_h^j defined as

$$\gamma_h^j = \frac{\alpha}{4\pi} Q_j^2 \left[-2 + \left(\frac{\alpha}{4\pi} \right) \frac{40}{9} \sum_k^{\text{massless}} n_k q_k^2 + \mathcal{O}(\alpha^3) \right]. \quad (5.60)$$

Finally, the third line of Eq. 5.56 contains the interference between initial-state and final-state emissions. We label with β_{j_1, j_2} the function defined by the relation

$$\cosh \beta_{j_1, j_2} = -\frac{p_{j_1} \cdot p_{j_2}}{m_{j_1} m_{j_2}}. \quad (5.61)$$

The physically allowed values for $\cosh \beta_{j_1, j_2}$ with j_1 and j_2 outgoing are $\cosh \beta_{j_1, j_2} \leq 1$, corresponding to $\beta_{j_1, j_2} = -b + i\pi$ with real $b \geq 0$.

The symbols Γ_n , Γ'_n in Eq. 5.54-5.55 indicate the coefficients of the series expansion in the weak coupling α of the corresponding anomalous dimensions

$$\Gamma^{(l)} = \sum_{n \geq 0} \left(\frac{\alpha}{4\pi} \right)^{n+1} \Gamma_n^{(l)}. \quad (5.62)$$

The sum in Eq. 5.55 takes care of the collinear emission of massive fermion pairs and of soft emission with a massive quark loop. As such, it runs over all the particles k with electric charge q_k considered with a non-zero mass m_k in the calculation.

The expression of $\mathcal{I}^{(0,2)}$ receives a contribution from the renormalization of the one-photon emission vertex, which has been performed in the $\overline{MS}$ scheme, and contributes with a β_0 term

$$\beta_0 = -\frac{4}{3} \sum_k^{\text{massless}} n_k q_k^2. \quad (5.63)$$

The expressions in Eqs. 5.54-5.55 are fully consistent with those presented in [276], and have been derived with the Abelianization of the QCD subtraction operator presented in [249, 255, 277].

5.3.2 The process

The process under study is the production of a pair of massive leptons in quark-antiquark annihilation, for the sake of definiteness an up-quark pair annihilating into an electron-positron pair

$$u(p_1) + \bar{u}(p_2) \rightarrow e^+(p_3) + e^-(p_4). \quad (5.64)$$




Figure 5.12: Sample two-loop Feynman diagrams: 1-particle irreducible and factorisable one-loop $\times$ one-loop contributions.

The electron mass is labeled by m_e . To regulate final state radiation, we consider the electron to be massive:

$$p_1^2 = p_2^2 = 0; \quad p_3^2 = p_4^2 = m_e^2. \quad (5.65)$$

Here we plan to study the IR structure of the second order purely EW corrections to this process, and to achieve this goal we limit the discussion to the evaluation in QED of the interference terms¹³

$$\langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(0,1)} \rangle, \quad \langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(0,2)} \rangle, \quad (5.66)$$

which contribute to the corrections to the unpolarized squared matrix element of the process in Eq. 5.64, at $\mathcal{O}(\alpha^2)$. The evaluation of the finite part of the two-loop virtual amplitude entails first the discussion of the UV renormalization, followed by the subtraction of the IR divergences.

In Figure 5.12 are presented some examples of two-loop Feynman diagrams contributing to the bare interference term $\langle \hat{\mathcal{M}}^{(0,0)} | \hat{\mathcal{M}}^{(0,2)} \rangle$. (a), (b), (c) are sample diagrams for each 1-particle irreducible (1PI) category in which we can group the diagrams, respectively self-energies, vertices and boxes 1PI diagrams. In contrast, (d), (e), (f) represent sample diagrams for the factorizable contributions given by the iteration of 2 one-loop diagrammatic corrections. The cancellation of the UV divergences stemming from the bare diagrams follows according to a standard renormalization procedure. The resulting amplitudes are expressed in terms of renormalized parameters, in turn related to measurable quantities. We write the UV-renormalized amplitude as a Laurent expansion in ϵ , and we keep only terms up to $\mathcal{O}(\epsilon^2)$ or $\mathcal{O}(\epsilon^0)$, at one- or two-loop respectively, in order to include all the second-order finite contributions when we perform the IR subtraction described in Section 5.3.1.

¹³We do not discuss the interference term $\langle \hat{\mathcal{M}}^{(0,1)} | \hat{\mathcal{M}}^{(0,1)} \rangle$, also contributing at $\mathcal{O}(\alpha^2)$, which can be easily obtained with any NLO amplitude generator.




Figure 5.13: Sample Feynman diagrams with the insertion of a closed fermionic loop along a photon line and the associated electric charge counterterm diagrams.

5.3.3 Ultraviolet renormalization

The SM renormalization up to second order in the electroweak interaction has been discussed in Sec. 2.1. The analytical expressions of the fermion mass and wave-function counterterms in QED for the case of one internal mass can be found in [276]. Our renormalization procedure within OCEANN is realized with a dedicated semi-analytical evaluation of the coupling and fermions renormalization constants, whose definitions follow from the 2-point Green functions of the theory and can be found in Appendix A. The generalization of this computational approach to the complete two-loop EW case is straightforward.

We write the virtual amplitude for the production of a pair of massive electrons as

$$|\hat{\mathcal{M}}(u\bar{u} \rightarrow e^+e^-)\rangle = 4\pi\hat{\alpha} \sum_{j=0}^{\infty} \left(\frac{\hat{\alpha}}{4\pi}\right)^j |\hat{\mathcal{M}}^{(0,j)}(s, t, \hat{m}_e, \hat{m}_f)\rangle, \quad (5.67)$$

where $|\hat{\mathcal{M}}^{(0,j)}\rangle$ is the bare amplitude at j^{th} perturbative order in $\hat{\alpha}$, and all the parameters (the electric charge $\hat{e}$, the electron mass $\hat{m}_e$ and the mass $\hat{m}_f$ of a generic fermion f) and fields are treated as bare quantities. The renormalized amplitude is obtained by the replacement of all the bare quantities with their renormalized counterparts: $\hat{e} \rightarrow e + \delta e$, $\hat{m}_e \rightarrow m_e + \delta m_e$, $\hat{m}_f \rightarrow m_f + \delta m_f$, and the rescaling $\hat{\psi} \rightarrow Z_{\psi}^{\frac{1}{2}}\psi = (1 + \delta Z_{\psi})^{\frac{1}{2}}\psi$ for each external fermionic field, where we consider $\hat{\alpha} = \hat{e}^2/(4\pi)$ and $\alpha = e^2/(4\pi)$,

$$\begin{aligned} |\mathcal{M}(u\bar{u} \rightarrow e^+e^-)\rangle &= Z_u^{\frac{1}{2}} Z_{\bar{u}}^{\frac{1}{2}} Z_{e^+}^{\frac{1}{2}} Z_{e^-}^{\frac{1}{2}} 16\pi^2 \times \\ &\quad \sum_{j=0}^{\infty} \left(\frac{e + \delta e}{4\pi}\right)^{2j+2} |\hat{\mathcal{M}}^{(0,j)}(s, t, m_e + \delta m_e, m_f + \delta m_f)\rangle \\ &= 4\pi\alpha \sum_{j=0}^{\infty} \left(\frac{\alpha}{4\pi}\right)^j |\mathcal{M}^{(0,j)}(s, t, m_e, m_f)\rangle. \end{aligned} \quad (5.68)$$

Since not only the bare amplitude, but also the counterterms are computed in perturbation theory, we expand the renormalized amplitude and collect all the terms of the same order in the coupling constant, checking the cancellation of the UV divergences. The result can be eventually cast as an expansion in the renormalized coupling α , where $|\mathcal{M}^{(0,j)}\rangle$ includes diagrammatic and counterterm contributions and is UV finite.

At first order, we identify three subsets of Feynman diagrams, the corrections to the tree-level photon propagator, the initial (final) vertex corrections and the associated wave function (WF) factors on the external quark (electron) lines, the box diagrams. Each of these three groups of diagrams is UV finite, after the inclusion of the one-loop counterterms: the boxes are UV finite by power counting, the photon

vacuum polarization is by construction made finite by the electric charge counterterm; the Ward identities guarantee the UV finiteness of the combination of the vertex corrections with the WF constants of the corresponding external lines. Additional poles in the dimensional regularization parameter ϵ are however still present, because of the IR divergent behaviour of the vertex and box corrections.

At the second perturbative order, the finiteness of the self-energy corrections to the tree-level photon propagator follows from the definition of the electric charge counterterm, once all the factorisable one-loop contributions have been properly included. Like in the one-loop case, the UV-finiteness of the sum of vertex and WF contributions can not be directly checked, because of the presence of IR poles, also expressed with the regulator ϵ . We identify two distinct groups of second-order 1PI diagrams: those with the exchange of two internal photons (like for instance diagrams (b) and (c) in Figure 5.12), and those with the insertion, along an internal photon line of a one-loop diagram, of a closed fermionic loop (like for instance diagram (a) in Figure 5.13). The latter features an UV divergence which is subtracted by the one-loop electric charge counterterm applied to the corresponding one-loop amplitude (diagrams (b) and (c) in Figure 5.13).

Since the IR subtraction terms discussed in Section 5.3.1 have been derived in the $\overline{MS}$ renormalisation scheme for the electric charge, we have to adopt it also in the evaluation of the virtual corrections to ensure a consistent cancellation of the IR divergencies. While our final result will thus be expressed in terms of the coupling $\alpha_{\overline{MS}}(\mu_R)$, evaluated at the renormalisation scale μ_R , it can also be expressed in terms of the on-shell fine structure constant $\alpha(0)$, with a finite renormalisation at the end of the calculation, as it is discussed in detail in Section 5.3.5. The fermion masses and the fermion wave-function counterterms are instead always computed in the on-shell scheme.

We remark that a change in the electric charge renormalisation scheme has a non trivial impact on the pattern of the subtraction of the IR poles. If we consider as an example the two diagrams (b) and (c) in Figure 5.13, we can easily recognize the fact that a finite additional contribution to the counterterm insertion is a constant which multiplies the entire underlying one-loop diagram. The latter features soft and collinear IR poles, which are multiplied by this additional factor. The cancellation of these extra poles is possible only with a corresponding contribution from the real emission diagrams, or equivalently from the subtraction term.

5.3.4 Explicit QED results and control over the cancellation of the IR poles

The calculation presented here illustrates with a non trivial example the application of the OCEANN framework presented in 4.4.1.

In the following, we present our results for the squared two-loop QED virtual amplitudes $\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,2)} \rangle$ for the process $u\bar{u} \rightarrow e^+e^-$. We identify in total 11 families relevant for the description of two-loop box topologies¹⁴. In the classification of the vertex corrections we identify 6 new families, while in the analysis of the self-energy diagrams we find 2 new families.

The fully diagrammatic approach allows us to perform intermediate checks of the calculation. We find a finite result for the self-energy corrections to the tree-level photon propagator, after the introduction of the one- and two-loop electric charge counterterms. We check the cancellation of one pole in ϵ , when we combine a two-loop diagram with one fermion loop insertion together with the charge renormalisation of the adjacent vertices in the corresponding one-loop diagram.

The UV-renormalised results described in the previous Section still feature IR divergences, which appear as poles in the dimensional regularization parameter ϵ , but also as logarithms of the lepton mass

¹⁴We count only once the insertion of a closed fermionic loop, treating the fermion mass as one generic parameter.

ϵ^{-4}	0.438 957 476 012 852 250 408 920 432
	-0.438 957 476 012 852 250 408 920 432
	0
ϵ^{-3}	-50.507 936 715 422 813 794 673 578 353
	50.507 936 715 422 813 794 673 578 353
	0
ϵ^{-2}	729.673 397 464 406 823 805 401 840 608
	-1334.179 413 878 226 352 371 754 281 425
	604.506 016 413 819 528 566 352 440 817
ϵ^{-1}	26 528.662 344 660 489 566 812 058 589 673
	-463.090 234 187 259 379 447 239 209 073
	-26 065.572 110 473 230 187 364 819 380 599
ϵ^0	401 847.447 345 130 812 729 654 039 382 635
	4376.812 767 994 162 824 254 022 583 940
	-308 980.388 535 504 055 198 013 748 925 268
$2\text{Re}\langle\mathcal{M}^{(0,0)} \mathcal{M}^{(0,2),fin}\rangle$	97 243.871 577 620 920 355 894 313 041 307

Table 5.1: Numerical coefficient, computed with at least 32 significant digits, of the IR poles present in the different contributions to Equation 5.53, for $s = 10000 \text{ GeV}^2$ and $t = -2500 \text{ GeV}^2$. For each pole, the first row refers to $2\text{Re}\langle\mathcal{M}^{(0,0)}|\mathcal{M}^{(0,2)}\rangle$, the second one to $-\mathcal{I}^{(0,2)} 2\text{Re}\langle\mathcal{M}^{(0,0)}|\mathcal{M}^{(0,0)}\rangle$ and the third one to $-\mathcal{I}^{(0,1)} 2\text{Re}\langle\mathcal{M}^{(0,0)}|\mathcal{M}^{(0,1),fin}\rangle$. The last row is the final result of $2\text{Re}\langle\mathcal{M}^{(0,0)}|\mathcal{M}^{(0,2),fin}\rangle$.

in the case of the final state collinear divergences. It is a Laurent series in the dimensional regularization parameter, and the coefficients of the various powers of ϵ are numbers with arbitrary precision. In Table 5.1 we present the value of the coefficients of the different powers of ϵ , for the diagrammatic part and for the two contributions to the subtraction operator, factorising the one-loop and Born squared amplitudes respectively. These results are obtained using the following values for the masses and Mandelstam invariants:

$s = 10000 \text{ GeV}^2$	$m_e = 0.510998 \text{ MeV}$
$t = -2500 \text{ GeV}^2$	$m_\mu = 0.10566 \text{ GeV}$
$\mu_R = m_Z = 91.1535 \text{ GeV}$	$m_\tau = 1.77686 \text{ GeV}$
	$m_t = 172.69 \text{ GeV}$

In the evaluation of the individual contributions to the amplitude, we encounter some terms whose value is several orders of magnitude larger than 1; we can inspect them all in absolute value and define $Q = \log_{10}(x)$, with x the largest term of the calculation. We consider verified the cancellation of the IR poles if, for any given number P of significant digits requested in the arbitrary precision evaluation framework, with $P > Q$, the sum of the coefficients yields 10^{-P+Q} , which is our explicit representation of zero. The results in Table 5.1 have been obtained with a precision request $P = 50$, while we observe a precision loss due to internal cancellations by $Q = 15$ orders of magnitude, eventually achieving control over 35 digits. The IR cancellation can be verified in greater detail by isolating subsets of contributions proportional to the same power of the electric charges of the initial state quark, final state lepton, and of the fermion running in the internal loops; each subset is separately IR finite.

Whenever possible, the precision losses can be reduced with a convenient choice of the MIs basis. In this computation we use the basis of Master Integrals automatically generated by the combined application of FEYNARTS, ABISS and KIRA, in order to test the efficiency of our code also with a suboptimal




Figure 5.14: We plot the contribution to $\langle \mathcal{M} | \mathcal{M} \rangle$ coming from one- and two-loop diagrams for different values of $\sqrt{s}$ and $t = -s/4$. In particular, in solid blue we present $(4\pi\alpha_{\overline{\text{MS}}}^3) 2\text{Re}\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,1),fin} \rangle$, while in dashed orange $\alpha_{\overline{\text{MS}}}^4 2\text{Re}\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,2),fin} \rangle$, with $\alpha_{\overline{\text{MS}}} = 1/128$. The two-loop contributions have been multiplied by a factor 10 to improve the readability.

choice of Master Integrals. The numerical control over the pole cancellation is the same as the one obtained using closed form expressions in terms of special functions for the Feynman integrals, given the possibility to push the precision to an arbitrary level. The solution of the MIs with a series expansion approach offers on the other hand a significant advantage for all those cases where a closed form solution is not viable: IR-finite integrals with several internal massive lines, but also IR-divergent integrals with 2 or 3 internal massive lines, which require a non trivial dedicated effort to be expressed in closed form [139].

In Table 5.1 we present also the result for the finite part of the interference term, which is in perfect agreement with the values presented in [276]. In Figure 5.14, we present additional values for the finite part of the interference term, calculated for different values of $\sqrt{s}$. The observed behavior explicitly illustrates the non-trivial interplay of the different powers of Sudakov logarithms present in the virtual corrections. Specifically, at moderate center-of-mass energies, the corrections initially grow due to the positive contribution of single-logarithmic terms. However, as $\sqrt{s}$ increases further, the negative double-logarithmic contributions gradually take over and dominate the asymptotic regime, eventually causing the overall corrections to decrease.

5.3.5 Impact of a change of scheme

The results discussed in Section 5.3.2 are computed in terms of the $\overline{\text{MS}}$ -renormalised electric charge $\alpha_{\overline{\text{MS}}}(\mu_R)$. The UV renormalised and IR subtracted squared amplitude, including the corrections up to

second order, reads:

$$\begin{aligned}
\langle \mathcal{M} | \mathcal{M} \rangle &= 16\pi^2 \alpha_{\overline{MS}}^2(\mu_R) \langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,0)} \rangle \\
&+ 4\pi \alpha_{\overline{MS}}^3(\mu_R) \left[2\text{Re} \left(\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,1),fin} \rangle \right) \right] \\
&+ \alpha_{\overline{MS}}^4(\mu_R) \left[2\text{Re} \left(\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,2),fin} \rangle \right) \right. \\
&\quad \left. + \langle \mathcal{M}^{(0,1),fin} | \mathcal{M}^{(0,1),fin} \rangle \right] + \dots \\
&= \alpha_{\overline{MS}}^2(\mu_R) \bar{\mathcal{A}}_{00} + \alpha_{\overline{MS}}^3(\mu_R) \bar{\mathcal{A}}_{01} + \alpha_{\overline{MS}}^4(\mu_R) [\bar{\mathcal{A}}_{02} + \bar{\mathcal{A}}_{11}] + \dots,
\end{aligned} \tag{5.69}$$

where the amplitudes $|\mathcal{M}^{(0,1),fin}\rangle$, $|\mathcal{M}^{(0,2),fin}\rangle$ and the abbreviations $\bar{\mathcal{A}}_{ij}$ depend on the renormalisation scale μ_R . Another choice to express the renormalised coupling is given by the on-shell scheme definition, at zero momentum transfer, $\alpha(0)$, from the study of the Thomson scattering. This second alternative provides an accurate description of the real-photon emission spectra, and it is thus of high phenomenological relevance. It is interesting to re-express Equation (5.69) in terms of the fine structure constant $\alpha(0)$, via a finite renormalisation. The relation between the on-shell and $\overline{MS}$ renormalised electric charges is given by [79]:

$$\alpha_{\overline{MS}}(\mu_R) = \frac{\alpha(0)}{1 - \Delta\alpha_{\overline{MS}}(\mu_R^2)}. \tag{5.70}$$

The definitions of the coefficients $\Delta\alpha_{\overline{MS}}^{(i)}(\mu_R^2)$ are reported in Appendix B, with the first order containing the hadronic part $\Delta\alpha_{\text{had}}(\mu_R^2)$, necessary to express the light-quarks contribution. This quantity is extracted from the experimental data of the cross section $e^+e^- \rightarrow \text{hadrons}$ and, in the case of $\mu_R = m_Z$, it is possible to use the determination from [278]: $\Delta\alpha_{\text{had}}(m_Z^2) = 0.027572 \pm 0.000359$. In the following we perform a technical comparison to test the perturbative convergence of the results when increasing the order of the perturbative approximation. Expanding Equation (5.70) one gets:

$$\begin{aligned}
\alpha_{\overline{MS}}(m_Z) &= \alpha(0) \left[1 + \alpha(0) \Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2) + \alpha(0)^2 \left(\left(\Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2) \right)^2 \right. \right. \\
&\quad \left. \left. + \Delta\alpha_{\overline{MS}}^{(2)}(m_Z^2) \right) \right].
\end{aligned} \tag{5.71}$$

Replacing the couplings in Equation (5.69) with Equation (5.71), we obtain:

$$\begin{aligned}
\langle \mathcal{M} | \mathcal{M} \rangle &= \alpha(0)^2 \bar{\mathcal{A}}_{00} + \alpha(0)^3 \left[\bar{\mathcal{A}}_{01} + 2\Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2) \bar{\mathcal{A}}_{00} \right] \\
&+ \alpha(0)^4 \left[\bar{\mathcal{A}}_{02} + \bar{\mathcal{A}}_{11} + 3\Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2) \bar{\mathcal{A}}_{01} \right. \\
&\quad \left. + 2\Delta\alpha_{\overline{MS}}^{(2)}(m_Z^2) \bar{\mathcal{A}}_{00} + 3 \left(\Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2) \right)^2 \bar{\mathcal{A}}_{00} \right] + \dots
\end{aligned} \tag{5.72}$$

In Table 5.2 we compare the sum of the contributions to the squared matrix element, using as renormalised electric charge $\alpha_{\overline{MS}}(m_Z)$ or $\alpha(0)$ ¹⁵. We consider different approximations: tree-level only, tree-level and one-loop corrections, the sum of tree-level, one-loop and of the two-loop $\langle \mathcal{M}^{(0,0)} | \mathcal{M}^{(0,2),fin} \rangle$ term. At tree level the difference between the results is purely due to the ratio of the coupling constants $(\alpha_{\overline{MS}}(m_Z)/\alpha(0))^2$. In the one-loop virtual corrections there are two subsets: the self-energy corrections

¹⁵Note that in order to use the formulae in [79], we have to account for the different normalisation of the Feynman integrals by a factor S_ϵ .

Highest order	$\alpha(0)$ scheme	$\alpha_{\overline{MS}}(m_Z)$ scheme	Δ (%)
tree-level	46.742	53.546	12.707
one-loop	86.610	94.297	8.152
two-loop	97.158	97.919	0.777

Table 5.2: Numerical results of the squared virtual amplitude at $\mathcal{O}(\alpha^2)$, $\mathcal{O}(\alpha^3)$ and $\mathcal{O}(\alpha^4)$ in the two schemes for the renormalised electric charge, with $\alpha_{\overline{MS}}(m_Z) = 1/128$ and $\alpha(0) = 1/137$, in units of 10^{-4} . The last column shows the percentage difference $\Delta = 100 |\langle \mathcal{M} | \mathcal{M} \rangle_0 / \langle \mathcal{M} | \mathcal{M} \rangle_{\overline{MS}} - 1|$.

together with the electric charge counterterms and the group composed by vertex and box corrections. The former, combined with the tree level, features in the two schemes much closer results compared to the tree-level case. The latter on the other hand are still evaluated with tree level couplings and, given the large difference between the couplings in the two schemes, they yield a difference between the results at the 8% level. It is thus interesting to observe the impact of the two-loop corrections, which include the renormalisation of the couplings present in the one-loop diagrams. Thanks to this additional step, also the one-loop vertices and boxes are evaluated with effective couplings closer in value, reducing the distance between the two schemes at sub percent level. The large difference between one perturbative order and the next one is due to the presence of double mass logarithms, which cancel upon inclusion of the real-emission contributions to the inclusive cross section [279].

5.4 Full Two-loop EW corrections: prospects

The computation of the full two-loop EW amplitude for $2 \rightarrow 2$ scattering processes represents a major milestone in the precision physics program, both for the High-Luminosity LHC and for future lepton colliders such as the FCC-ee. Achieving this level of precision is essential to match the anticipated sub-percent experimental uncertainties in crucial measurements, such as the W boson mass and the effective weak mixing angle.

While some partial results already exist for specific channels such as $e^+e^- \rightarrow Zh$, where the gauge-invariant subset of contributions involving closed fermion loops was computed in [280], a complete phenomenological calculation remains elusive. Recently, the full bare two-loop amplitude for this process was presented in [281]. However, it lacks UV renormalization and IR subtraction, rendering it not yet suitable for a full NNLO cross-section calculation. Concerning the DY process, while the QED contributions at $\mathcal{O}(\alpha^2)$ have been discussed in Section 5.3, the full two-loop EW amplitudes are not yet available in the literature.

5.4.1 Diagrams

The first computational challenge encountered when evaluating full two-loop EW contributions to DY is the sheer proliferation of Feynman diagrams. This dramatic scaling is driven by the high multiplicity of particles in the EW sector, including multiple massive gauge bosons, the Higgs boson, and ghosts.

Table 5.3 explicitly details the number of diagrams that must be evaluated at different perturbative orders. The table highlights not only the rapid growth associated with higher loop orders but also the huge difference in complexity between purely QCD and purely EW corrections. This huge number of diagrams necessitates fully automated handling, for instance, by utilizing the OCEANN framework presented in Section 4.4.1. Furthermore, the analytical expressions for individual diagrams become

	NC-DY	CC-DY
tree-level	2	1
one-loop QCD	2	1
one-loop EW	101	44
two-loop QCD	58	17
two-loop QCD-EW	1428	209
two-loop EW	9019	4536

Table 5.3: Number of Feynman diagrams that need to be considered at different perturbative orders.




Figure 5.15: Examples of two-loop boxes which contribute to the two-loop EW corrections for the NC-DY.

exceedingly complex at higher perturbative orders due to the increased number of internal lines and distinct mass scales. This leads to substantially larger analytical expressions and severely increased processing times for evaluating the interference terms. As an example, computing the interferences between the Born amplitude and the $\mathcal{O}(10^3)$ two-loop boxes at $\mathcal{O}(\alpha^2)$ for the NC-DY process required $\mathcal{O}(8h)$ of CPU time using ABISS, yielding an output file of $\mathcal{O}(120 \text{ MB})$. A sample of these diagram topologies is illustrated in Fig. 5.15.

5.4.2 Integrals and differential equations

The second major bottleneck lies in the IBP reduction of Feynman integrals and the subsequent evaluation of the MIs. This difficulty is directly tied to the multitude of internal massive lines characteristic of the EW sector. For instance, the second diagram in Fig. 5.15 contains five massive lines characterized by two distinct mass scales. Combining these masses with the two Mandelstam invariants and the dimensional regulator ϵ , the final integrals depend on up to five variables.

This multi-scale dependence makes the reduction to MIs highly demanding in terms of both CPU time and memory. In our studies, we found that reducing these topologies often exceeded the baseline capabilities of KIRA 3.0. However, by coupling KIRA with the rational arithmetic tracker RATRACER, we successfully reduced all scalar integrals to their respective MIs. Specifically, reducing the most complicated two-loop box topologies required running on a high-performance cluster using $\mathcal{O}(100)$ cores for roughly 50 h, and using approximately 150 GB of RAM.

Reducing the integrals to a basis of MIs is only the first step. The next phase is their numerical evaluation, which is typically achieved by deriving a system of differential equations, calculating the boundary conditions at a specific point in the physical phase space, and finally solving the system to generate a numerical grid suitable for Monte Carlo generators. Each of these steps introduces unique challenges.

Concerning the evaluation of the boundary conditions, we employed `AMFLOW` to evaluate them at a specific phase-space point. However, for some of the most intricate topologies, the program failed to converge. This instability arises because our calculation utilizes complex masses for W and Z bosons. Internally, `AMFLOW` relies on `KIRA` to construct the system with respect to the auxiliary mass parameter η . Because `KIRA` does not natively support fixing kinematic parameters to complex floating-point values, the IBP reduction must retain the exact analytical dependence on the complex mass. This drastically slows down, and in extreme cases completely halts, the generation of the DEs. Given that our standard reductions succeeded when pairing `KIRA` with `RATRACER`, integrating a similar rational reconstruction approach directly within `AMFLOW` could potentially solve this issue.

Generating the DE system for subsequent integration via `SEASYDE` requires keeping the size of the coefficient matrix manageable, as the system size directly dictates the evaluation runtime. In the mixed QCD-EW calculation for the NC-DY process, the DE system for the most complex topology was approximately 700 kB, and the grid generation took about 5 days. For the CC-DY process, the largest topology reached 2.1 MB, requiring roughly 10 days of `SEASYDE` runtime. However, for the full two-loop EW calculation, the size of some systems explodes to $\mathcal{O}(350 \text{ MB})$, making direct integration practically unfeasible. To mitigate the issue of massive DE systems, two primary strategies can be employed. On the technical side, the rational reconstruction techniques implemented in packages like `FINITEFLOW` are highly beneficial. In particular, by simplifying the algebraic expressions of the matrix coefficients, one can achieve a file size reduction of up to 50%. On the physical side, transitioning to a different basis of MIs, such as a canonical basis, would dramatically simplify the DE system. However, discovering such a basis for multi-scale problems is highly non-trivial. Automated procedures are currently limited, generally requiring a significant amount of manual intervention and heuristic approaches.

As detailed in Table 5.4, the benchmark numbers explicitly illustrate the escalation in computational complexity when transitioning from mixed QCD-EW to full two-loop EW corrections for the DY processes. These specific topologies were chosen as a representative sample, ordered by increasing difficulty, to showcase the varying levels of complexity one must face in these calculations. The data highlights how each stage of the calculation is stressed by the inclusion of the full EW sector.

The first major indicator of complexity is the number of MIs. For the mixed two-loop NC-DY and CC-DY topologies, the number of MIs is relatively modest (36 and 56, respectively). Consequently, the IBP reduction using `KIRA` combined with `FIREFLY` is manageable, taking 12 to 16 hours on a standard 32-core node. However, as we move to the full two-loop EW NC-DY topologies, the number of masters increases significantly, peaking at 140 for the most intricate crossed-box topology. The reduction time for this specific topology climbs to 54 hours, even when utilizing a much larger 120-core cluster and integrating the `RATRACER` package to optimize the reduction procedure.

The evaluation of BCs at a single phase-space point using `AMFLOW` demonstrates a similar scaling. While the mixed topologies require roughly an hour to evaluate, the full EW topologies demand considerably more resources, taking up to 4 hours on 120 cores for the 126 master topology. Notably, for the 140 master topology, `AMFLOW` was entirely unable to compute the BCs. As previously discussed, this failure is a direct consequence of using the complex mass.

Finally, the size of the systems of differential equation provides a clear metric for the memory and integration bottlenecks encountered during the final stage. The mixed NC-DY system is a highly manageable 700 kB, allowing `SEASYDE` to successfully generate a 3250 point numerical grid in 5 days. The mixed CC-DY system grows to 2.1 Mb, which correspondingly doubles the `SEASYDE` runtime to 10 days. In sharp contrast, the full two-loop EW systems explode in size, reaching up to 45 Mb and 350 Mb for the more complex boxes. Consequently, the grid generation times for these topologies are currently

	NC-DY 2L mixed	CC-DY 2L mixed	Z on-shell 2L EW	NC-DY 2L EW	NC-DY 2L EW	NC-DY 2L EW	NC-DY 2L EW
Topology							
# MIs	36	56	51	47	104	126	140
IBPs KIRA +FIREFLY	12 h 32 cores	16 h 32 cores	1 d 32 cores	1.5 m 96 cores + RATRACER	30 m 120 cores +RATRACER	8 h 120 cores +RATRACER	54 h 120 cores +RATRACER
BCs AMFLOW	50 m 32 cores	75 m 32 cores	6.75 h 32 cores	15 m 96 cores	1 h 120 cores	4 h 120 cores	Did not converge
Size DEs	700 Kb	2.1 Mb	NA	4 Mb	45 Mb	350 Mb	??
3250 pt grid SEASYDE	5 d	10 d	NA	??	??	??	??

Table 5.4: In this Table we list the reduction times and dimensions of the final system of equations for various benchmark topologies. In the diagrams, straight lines represent massless fermions and curly lines denote gluons. The blue dashed line corresponds to the Higgs boson. Wavy lines indicate vector bosons: black for the photon, red for the Z boson, and green for the W boson.

undetermined, as systems of this magnitude are computationally prohibitive to integrate directly without further algebraic simplification or a transformation to a more optimal basis of MIs.

Beyond the generation of the numerical grids and the evaluation of the MIs, another severe numerical bottleneck is the final assembly of the amplitude. When combining the numerically evaluated MIs with their respective rational coefficients, huge numerical cancellations occur. These cancellations are driven by the intricate algebraic structure of the reduction and the spurious kinematic singularities present in the coefficients. For instance, in our calculations for both the NC- and CC-DY, we evaluated the necessary MIs with 20 significant digits. However, after combining them with their coefficients, we obtained a final result accurate to just 8 significant digits. This dramatic loss of precision implies that the MIs have to be computed with a significantly inflated number of digits to guarantee a stable and phenomenologically useful final amplitude. Unfortunately, the exact target precision cannot be easily predicted in advance and is typically determined a posteriori by observing the stability of the final result. A potential mitigation to this issue is, once again, the transformation to a different basis of MIs, such as the canonical base. A carefully chosen basis, indeed, could feature rational coefficients with a different kinematic dependence, therefore reducing the numerical cancellations and the precision requirements imposed on the numerical integration stage. The systematic determination of transformations capable of casting these differential equation systems into a canonical, ϵ -factorized form remains a highly active and complex field of ongoing research [282].

Phenomenology

“The mountaineer is a man who leads his body where, one day, his eyes have looked.”

Gaston Rébuffat, *French alpinist*.

In this chapter we present a detailed study on the impact of mixed QCD-EW corrections on the neutral-current DY process at LHC energies. We start by considering bare muons and examine the mixed QCD-EW corrections in both the resonant region and in the region of high invariant masses. We then analyse the impact of radiative corrections on the forward-backward asymmetry. We finally consider the case of dressed leptons. Here we show that our calculation can be extended to the massless limit, and we present a comparison of our results with those of the massless calculation of [217]. This study has been presented in [4].

6.1 Computational details

We consider the inclusive production of a pair of massive muons in proton–proton collisions,

$$pp \rightarrow \mu^+ \mu^- + X. \quad (6.1)$$

The theoretical predictions for this process can be obtained as a convolution of the parton distribution functions for the incoming protons with the hard-scattering partonic cross sections. When QCD and EW radiative corrections are considered, the initial partons include quarks, anti-quarks, gluons and photons.

The differential cross section for the process in Eq. (6.1) can be written as

$$d\sigma = \sum_{m,n=0}^{\infty} d\sigma^{(m,n)}, \quad (6.2)$$

where $d\sigma^{(0,0)} \equiv d\sigma_{\text{LO}}$ is the Born level contribution and $d\sigma^{(m,n)}$ the $\mathcal{O}(\alpha_S^m \alpha^n)$ correction. The mixed QCD–EW corrections correspond to the term $m = n = 1$ in this expansion and include double-real, real–virtual and purely virtual contributions. The corresponding tree-level and one-loop scattering amplitudes are computed with OPENLOOPS [163, 283, 284] and RECOLA [164, 285, 286], the results being in complete agreement¹. The two-loop amplitude has been computed in Section 5.1. Even when all the amplitudes have been computed, the completion of the calculation remains a formidable task. Indeed, double-real, real–virtual and purely virtual contributions are separately infrared divergent, and a method to handle and cancel infrared singularities has to be worked out. In this calculation we use a formulation

¹To be precise, the results obtained with OPENLOOPS and RECOLA coincide pointwise in phase space when adopting the G_μ or $\alpha(M_Z)$ EW renormalization schemes. They differ in the case of the $\alpha(0)$ scheme because of the different treatment of the light-fermion contributions to the photon vacuum polarization.

of the q_T subtraction formalism [157] derived from the NNLO QCD computation of heavy-quark pair production [259–261] through an appropriate Abelianization procedure [203, 262]. According to the q_T subtraction formalism [157], $d\sigma^{(m,n)}$ can be evaluated as

$$d\sigma^{(m,n)} = \mathcal{H}^{(m,n)} \otimes d\sigma_{\text{LO}} + \left[d\sigma_{\text{R}}^{(m,n)} - d\sigma_{\text{CT}}^{(m,n)} \right]. \quad (6.3)$$

The first term in Eq. 6.3 is obtained through a convolution (denoted by the symbol $\otimes$) of the perturbatively computable function $\mathcal{H}^{(m,n)}$ and the LO cross section $d\sigma_{\text{LO}}$, with respect to the longitudinal-momentum fractions of the colliding partons. The second term is the *real* contribution $d\sigma_{\text{R}}^{(m,n)}$, where the charged leptons are accompanied by additional QCD and/or QED radiation that produces a recoil with finite transverse momentum q_T . For $m + n = 2$ such contribution can be evaluated by using the dipole subtraction formalism [155, 156, 246, 287–291]. In the limit $q_T \rightarrow 0$ the real contribution $d\sigma_{\text{R}}^{(m,n)}$ is divergent, since the recoiling radiation becomes soft and/or collinear to the initial-state partons. The counterterm $d\sigma_{\text{CT}}^{(m,n)}$ is explicitly designed to replicate and cancel this singular behavior, ultimately rendering the full cross-section in Eq. 6.3 finite. However, due to the non-local nature of the q_T -subtraction method, this exact cancellation cannot be evaluated directly at $q_T = 0$ point-by-point in a four-dimensional Monte Carlo integration without encountering severe numerical instabilities. Consequently, the numerical integration of the quantity in the square brackets of Eq. 6.3 is only possible when a finite kinematic cutoff is applied. In practice, the integration is restricted to the well-defined region $q_T/m_{\mu\mu} > r_{\text{cut}}$, where $m_{\mu\mu}$ is the invariant mass of the dimuon system. To recover the physical, cutoff-independent prediction, the numerical evaluation is performed for several progressively smaller values of r_{cut} , and the final result is extracted by carrying out an extrapolation to the $r_{\text{cut}} \rightarrow 0$ limit (see Ref. [162] for more details).

The required phase space generation and integration is carried out within the MATRIX framework [162]. The core of MATRIX is the Monte Carlo program MUNICH², which contains a fully automated implementation of the dipole subtraction method for massless and massive partons at NLO QCD [155, 156, 246] and NLO EW [287–291]. The q_T subtraction method has been applied to several NNLO QCD computations for the production of colourless final-state systems (see Ref. [162] and references therein), and to heavy-quark pair production [259–261] and related processes [292–295], which correspond to the case $m = 2, n = 0$. The method has also been used in Ref. [262] to study NLO EW corrections for the DY process, which represents the case $m = 0, n = 1$. The structure of the coefficients $\mathcal{H}^{(1,1)}$ and $d\sigma_{\text{CT}}^{(1,1)}$ can be derived from those controlling the NNLO QCD computation of heavy-quark pair production. The initial-state soft/collinear and purely collinear contributions were already presented in Ref. [204]. The fact that the final state is colour neutral implies that final-state radiation is of pure QED origin. Therefore, the purely soft contributions have a simpler structure than the corresponding contributions entering the NNLO QCD computation of heavy-quark pair production [296].

6.2 Phenomenological results

Unless stated otherwise, we work in the G_μ scheme with $G_F = 1.1663787 \times 10^{-5} \text{ GeV}^{-2}$ and set the *on-shell* values of masses and widths to $m_{W,\text{OS}} = 80.385 \text{ GeV}$, $m_{Z,\text{OS}} = 91.1876 \text{ GeV}$, $\Gamma_{W,\text{OS}} = 2.085 \text{ GeV}$, $\Gamma_{Z,\text{OS}} = 2.4952 \text{ GeV}$. Those values are translated to the corresponding *pole* values $m_V = m_{V,\text{OS}}/\sqrt{1 + \Gamma_{V,\text{OS}}^2/m_{V,\text{OS}}^2}$ and $\Gamma_V = \Gamma_{V,\text{OS}}/\sqrt{1 + \Gamma_{V,\text{OS}}^2/m_{V,\text{OS}}^2}$, $V = W, Z$, from which $\alpha = \sqrt{2} G_F m_W^2 (1 - m_W^2/m_Z^2)/\pi$ is derived, and we use the complex-mass scheme [98] throughout. The

²MUNICH, which is the abbreviation of MUlti-chaNnel Integrator at Swiss (CH) precision, is an automated parton-level NLO generator by S. Kallweit.

muon mass is fixed to $m_\mu = 105.658369$ MeV, and the pole masses of the top quark and the Higgs boson to $m_t = 173.07$ GeV and $m_H = 125.9$ GeV, respectively. The CKM matrix is taken to be diagonal. We work with $n_f = 5$ massless quark flavours and retain the exact top-mass dependence in all virtual and real–virtual amplitudes associated with bottom-induced processes, except for the two-loop virtual corrections, where we neglect top-mass effects. Given the smallness of the bottom-quark density, we estimate the corresponding error to be at the percent level of the computed correction. We use the NNPDF31_nnlo_as_0118_luxqed set of parton distributions [297], which is based on the LUXqed methodology [228] for the determination of the photon density. If not stated otherwise, the central values of the renormalization and factorization scales are fixed to $\mu_R = \mu_F = m_Z$ (and the corresponding value of α_S is set), while theoretical uncertainties are estimated by the customary 7-point scale variation, i.e. by varying μ_R and μ_F by a factor of two around their central values with the constraint $1/2 < \mu_R/\mu_F < 2$. We consider bare muons in the final state, unless stated otherwise.

Our results for the mixed QCD–EW corrections will be compared with those from an approach in which QCD and EW corrections are assumed to completely factorize. Such approximation can be defined as follows: for each bin of a distribution $d\sigma/dX$, the QCD correction, $d\sigma^{(1,0)}/dX$, and the EW correction restricted to the $q\bar{q}$ channel, $d\sigma_{q\bar{q}}^{(0,1)}/dX$, are computed, and the factorized $\mathcal{O}(\alpha_S\alpha)$ correction is calculated as

$$\frac{d\sigma_{\text{fact}}^{(1,1)}}{dX} = \left(\frac{d\sigma^{(1,0)}}{dX} \right) \times \left(\frac{d\sigma_{q\bar{q}}^{(0,1)}}{dX} \right) \times \left(\frac{d\sigma_{\text{LO}}}{dX} \right)^{-1}. \quad (6.4)$$

This approximation is justified if the dominant sources of QCD and EW corrections factorise with respect to the hard gauge-boson production subprocess (see the discussion in Ref. [215]).

6.2.1 Resonant region

We first study the impact of the mixed corrections on the bulk of the DY cross section around the Z peak. We consider proton–proton collisions at $\sqrt{s} = 13.6$ TeV and a fiducial volume defined by staggered cuts³ on the transverse momenta of the positively/negatively charged leptons,

$$p_{T,\mu^+} > 27 \text{ GeV}, \quad p_{T,\mu^-} > 25 \text{ GeV}, \quad |y_\mu| < 2.5, \quad (6.5)$$

while the invariant mass of the dimuon pair $m_{\mu\mu}$ fulfils

$$66 \text{ GeV} < m_{\mu\mu} < 116 \text{ GeV}. \quad (6.6)$$

In Table 6.1 we show our predictions for the fiducial cross section with an increasing radiative content. We start our discussion from the pure QCD predictions, labelled LO, NLO_{QCD} and NNLO_{QCD}. The NLO QCD corrections increase the fiducial cross section by about 7% with respect to the LO, their perturbative uncertainty, estimated through scale variations, being about 4%. The NNLO QCD corrections are negative, and amount to about –2%, while the associated uncertainties are reduced to below the percent level. N³LO corrections with this setup are not available, but their impact is likely to exceed the NNLO scale uncertainties [200].

Predictions including NLO EW effects are shown in the fourth and fifth rows of Table 6.1, and are labelled NLO_{EW} and NNLO_{QCD+NLO_{EW}}, respectively, while our best prediction, which includes mixed QCD–EW corrections as well, is labelled NNLO_{QCD+MIX+NLO_{EW}}. We see that NLO EW corrections

³This combination of cuts is known to restore the quadratic dependence of the acceptance in the transverse momentum of the dilepton pair, see e.g. Refs [162, 298].

are negative, about -4% , and have a larger impact than what might be expected by a naive coupling counting. Indeed, they are similar in size to the NLO QCD effects and, being negative, largely compensate the latter. We observe that the scale dependence of the NLO_{EW} prediction is identical to that at LO, as might have been expected, given that the relevant amplitudes do not involve the QCD coupling α_S . The newly computed mixed QCD–EW corrections are small for this particular setup, being below half a percent. They slightly reduce the fiducial cross sections. The small size of the mixed QCD–EW corrections is consistent with the fact that they do not reduce the scale dependence.

	σ [pb]	$\sigma^{(i,j)}$ [pb]	$\sigma^{(i,j)}/\sigma_{\text{LO}}$
LO	$735.73(3)^{+12.7\%}_{-13.6\%}$	—	—
NLO_{QCD}	$790.23(4)^{+2.7\%}_{-4.4\%}$	$\sigma^{(1,0)} = +54.50(5)$	$+7.4\%$
NNLO_{QCD}	$774.4(7)^{+0.7\%}_{-0.7\%}$	$\sigma^{(2,0)} = -15.8(5)$	-2.2%
NLO_{EW}	$704.07(3)^{+12.7\%}_{-13.6\%}$	$\sigma^{(0,1)} = -31.66(3)$	-4.3%
$\text{NNLO}_{\text{QCD}}+\text{NLO}_{\text{EW}}$	$742.7(5)^{+0.3\%}_{-0.7\%}$	—	—
$\text{NNLO}_{\text{QCD}}+\text{MIX}+\text{NLO}_{\text{EW}}$	$739.6(5)^{+0.6\%}_{-0.6\%}$	$\sigma^{(1,1)} = -3.1(2)$	-0.4%

Table 6.1: Fiducial cross section at the different perturbative orders and the corresponding corrections $\sigma^{(i,j)}$ as defined in Eq. (6.2).

We now turn to differential distributions. In Fig. 6.1 (left) we show our results for the rapidity distribution $y_{\mu\mu}$ of the dimuon pair. The strongest shape effect can be observed at NLO QCD and to a lesser extent at NNLO QCD. The inclusion of the mixed QCD–EW corrections further distorts the shape of the distribution up to 1% in the region $|y_{\mu\mu}| > 1.5$, and could be potentially relevant for a determination of the proton PDFs with 1% precision [299].

In Fig. 6.1 (right), we consider the invariant-mass distribution of the dimuon system. QCD radiative corrections are generally mild for this observable, while the emission of QED radiation off the final-state lepton pair has a significant impact if the invariant mass is reconstructed from bare leptons, as we do here. The inclusion of the NLO EW corrections indeed leads to a shift of events, giving rise to the formation of the characteristic radiative tail in the region on the left of the peak.⁴ Purely weak NLO corrections modify the strength of the Z -boson couplings to fermions and in turn, at LO, the squared matrix element of the Z -exchange diagram and the γ – Z interference. They affect the region around the Z resonance, with positive corrections reaching up to a few percent below the resonance and moderate negative corrections above the resonance [214].

Focusing on the impact of the mixed QCD–EW corrections in the middle plot, we observe sizeable effects and a non-trivial shape distortion. The correction is vanishing around the peak, where it reaches the maximum slope, while it is rather flat, positive and around 4% in the region below the peak and rather flat, negative and around -2% in the region above. We observe that the prediction obtained with the factorised ansatz (see Eq. (6.4)) works well in the latter region, while it fails to describe the exact result below the Z resonance, because of a non-trivial kinematical interplay between the production of the gauge boson and its subsequent decay into leptons, as it has been first pointed out in Ref. [213].

⁴The effect is also present but reduced by approximately a factor of two when leptons and photons are recombined [193].




Figure 6.1: Predictions for the rapidity distribution (left) and the invariant-mass distribution (right) of the final-state muon pair, in the setup defined in Section 6.2.1. NLO EW and mixed QCD–EW corrections are additively included on top of the NNLO QCD prediction. The relative effects of the mixed corrections on top of the additive combination NNLO QCD + NLO EW are shown in the middle panel, comparing the exact mixed corrections to the approximate factorised ansatz defined in the text. NLO EW as well as NLO and NNLO QCD K -factors are displayed in the bottom panel.

Furthermore, we notice that the inclusion of the mixed corrections leads to a significant reduction of the uncertainty band.

In Fig. 6.2 we show the rapidity distribution of the μ^+ (left) and the $\cos \theta^*$ distribution (right). The Collins–Soper (CS) angle θ^* [300] is defined in terms of LAB frame variables as

$$\cos \theta^* = \frac{y_{\mu\mu}}{|y_{\mu\mu}|} \frac{2(p_{\mu^-}^+ p_{\mu^+}^- - p_{\mu^-}^- p_{\mu^+}^+)}{m_{\mu\mu} \sqrt{m_{\mu\mu}^2 + p_{T,\mu\mu}^2}}, \quad (6.7)$$

with $p^\pm = (p^0 \pm p^3)/\sqrt{2}$. The factor $y_{\mu\mu}/|y_{\mu\mu}|$ takes into account the fact that in hadron collisions the quark and anti-quark directions are not known for each collision. However, on average the quark carries more energy than the anti-quark as the latter must originate from the sea. Then, the produced dilepton system is, on average, boosted in the direction of the valence quark. Since the Z boson rapidity is correlated with the direction of the valence quark, it offers a sensible reference to define the scattering angle.

The radiative corrections to the rapidity distribution of the anti-muon (Fig. 6.2 (left)) have a rather uniform impact over the considered rapidity range. In the case of the $\cos \theta^*$ distribution, radiative corrections distort the distribution in different ways. NLO QCD corrections are larger at large $|\cos \theta^*|$, which is caused by the qg channel opening up at NLO, dominated by configurations where a valence quark at relatively large momentum fraction x and a gluon at small x enter the hard process. By contrast, EW corrections distort the distribution in the opposite way, their impact being slightly larger and negative at




Figure 6.2: Same as Fig. 6.1 for the rapidity distribution of the anti-muon (left) and the CS angle $\cos \theta^*$ (right).

large values of $|\cos \theta^*|$. The impact of the mixed QCD–EW corrections is generally below the percent level, reaching the few percent level in the region of large $|\cos \theta^*|$.

6.2.2 High invariant-mass region

Besides the resonant region, percent-level precision can be envisaged also for data taken in the region of high invariant masses. This requires a careful assessment of the impact of EW radiative corrections and their interplay with QCD effects. Following Ref. [301], we consider proton–proton collisions at $\sqrt{s} = 13$ TeV and impose the following cuts on the final-state leptons:

$$p_{T,\mu^\pm} > 53 \text{ GeV}, \quad |y_\mu| < 2.4, \quad m_{\mu\mu} > 150 \text{ GeV}. \quad (6.8)$$

For this setup, a dynamical scale is a more appropriate choice. We use as central values of the renormalization and factorization scales the invariant mass of the dimuon system $m_{\mu\mu}$.

In Fig. 6.3 we show the invariant-mass distribution up to $m_{\mu\mu} = 2$ TeV. Excluding the first bin close to the edge at $m_{\mu\mu} = 150$ GeV, imposed by the selection requirements on the leptons, the effect of the mixed corrections is negative and increasing at higher invariant masses, reaching -5% at $m_{\mu\mu} = 2$ TeV. Such values are comparable with or even larger than the statistical error expected at the end of the High-Luminosity (HL) phase of the LHC, with a total collected luminosity of 3 ab^{-1} . The inclusion of these corrections should help reducing the theoretical uncertainty on the partonic cross section of the DY process and, in turn, help constraining the proton PDFs, at large partonic x , in a global fit of collider data.⁵ The precise knowledge of the higher-order corrections to this distribution is the mandatory starting point for an analysis which aims at the determination of the running of the $\overline{\text{MS}}$ weak mixing angle [303].

⁵The possibility that a PDF fit at large invariant masses inadvertently reabsorbs a New Physics signal has been discussed in Ref. [302].




Figure 6.3: Dimuon invariant-mass distribution, in the acceptance setup defined in Section 6.2.2. The structure of the panels and the color codes are the same as in Fig. 6.1.

We observe that the pattern of the mixed QCD–EW corrections is well approximated by the factorised ansatz (see Eq. (6.4)). This implies that QCD and EW effects largely factorise in this region of high invariant masses, as also observed in Ref. [217].

In Fig. 6.4, we show the invariant-mass distributions in the backward and forward regions up to $m_{\mu\mu} = 6$ TeV. We see that the impact of mixed corrections increases as $m_{\mu\mu}$ increases. The effect is slightly more pronounced in the forward region than in the backward region, following an analogous effect in the NLO EW corrections. As observed in Fig. 6.3 the factorised ansatz works rather well.

6.2.3 Forward–Backward asymmetry

The Forward–Backward (FB) asymmetry $A_{FB}(m_{\ell\ell})$ is one of the main observables for the determination of the leptonic effective weak mixing angle $\sin^2 \theta_{\text{eff}}^\ell$. At the LHC, $A_{FB}(m_{\ell\ell})$ is defined as the difference between the cross sections in the forward and backward directions normalized to the total cross section, differentially with respect to the invariant mass $m_{\ell\ell}$ of the dilepton system [191, 304],

$$A_{FB}(m_{\ell\ell}) = \frac{F(m_{\ell\ell}) - B(m_{\ell\ell})}{F(m_{\ell\ell}) + B(m_{\ell\ell})}. \quad (6.9)$$

The forward and backward cross sections are given by

$$\begin{aligned} F(m_{\ell\ell}) &= \int_0^1 d\cos\theta^* \frac{d\sigma}{d\cos\theta^* dm_{\ell\ell}}(m_{\ell\ell}), \\ B(m_{\ell\ell}) &= \int_{-1}^0 d\cos\theta^* \frac{d\sigma}{d\cos\theta^* dm_{\ell\ell}}(m_{\ell\ell}), \end{aligned} \quad (6.10)$$

where $\cos\theta^*$ is defined in Eq. (6.7).




Figure 6.4: Dimuon invariant-mass distribution, in the acceptance setup defined in Section 6.2.2, in the CMS bins [301], in the backward (left) and forward (right) regions. The structure of the panels and the colour codes are the same as in Fig. 6.1.

The FB asymmetry is defined from the identification of the scattering angle of the negatively charged outgoing lepton, with respect to the direction of the incoming particle, in our case at LO a quark. In proton–proton collisions, we cannot prepare the partonic initial state or, in other words, we are not aware of the direction of the incoming quark: each final state receives, at LO, contributions from both quark–antiquark and antiquark–quark annihilation, listing the partons involved in the hard scattering process from the first and the second hadron, respectively. If the partonic centre-of-mass system is at rest in the laboratory frame, $y_{\ell\ell} = 0$, the invariance of the system under interchange of the incoming hadrons exactly cancels out all the parity-violating contributions to $A_{FB}(m_{\ell\ell})$. At non-vanishing $y_{\ell\ell}$ values, the quark–antiquark subprocess (with the incoming quark oriented along the rapidity direction) prevails over the antiquark–quark process, because of the larger weight of the quark PDF with its valence content than that of the antiquark PDF; the unbalance avoids the cancellation of the parity-violating effects, and we observe a non-vanishing $A_{FB}(m_{\ell\ell})$. The slope of the asymmetry, in the Z resonance region, is steeper for larger $y_{\ell\ell}$ values, as illustrated by the different panels in Fig. 6.5.

The experiments run at LEP/SLC [9] have measured the leptonic FB asymmetry at the Z resonance with a precision of about 6%, which in turn allowed a determination of $\sin^2 \theta_{\text{eff}}^\ell$ with a relative error of about $7 \cdot 10^{-4}$. Measurements of the FB asymmetry have been performed by the CMS experiment at 8 TeV [305], by the ATLAS experiment at 7 TeV [306] and by the LHCb experiment exploiting data at 7 and 8 TeV [307]. More recently, the CMS experiment has measured $A_{FB}(m_{\ell\ell})$ and extracted $\sin^2 \theta_{\text{eff}}^\ell$ from data collected during Run 2 at 13 TeV [308]. These measurements have already reached a precision better than the percent level.

Given the foreseen increase in statistics in the HL-LHC phase, it is important to investigate the impact of higher-order QCD, EW, and mixed QCD–EW radiative corrections in the theory predictions of $A_{FB}(m_{\ell\ell})$. These effects have been recently studied in the pole approximation [214] in proton–proton




Figure 6.5: Theoretical predictions for the FB asymmetry for different slices of rapidity of the dimuon final state at different perturbative orders. Shifts with respect to the LO prediction induced by separate sets of corrections are reported in the bottom panels.

collisions at 13 TeV. In the following, we present a study on the impact of the exact mixed QCD–EW corrections to the FB asymmetry, for dimuon production in proton–proton collisions at 8 TeV. We adopt the setup used by the CMS collaboration for the extraction of $\sin^2 \theta_{\text{eff}}^\ell$ [305]:

$$p_{T,1} > 25 \text{ GeV}, \quad p_{T,2} > 15 \text{ GeV}, \quad |y_\mu| < 2.4, \quad 60 < m_{\mu\mu} < 120 \text{ GeV}, \quad (6.11)$$

where $p_{T,1}$ ($p_{T,2}$) is the transverse momentum of the hardest (second hardest) lepton. We consider five intervals in the absolute rapidity of the dimuon final state: $0 < |y_{\mu\mu}| < 0.4$, $0.4 < |y_{\mu\mu}| < 0.8$, $0.8 < |y_{\mu\mu}| < 1.2$, $1.2 < |y_{\mu\mu}| < 1.6$ and $1.6 < |y_{\mu\mu}| < 2.4$. In order to better quantify the impact of the higher-order corrections, we consider the shifts with respect to the LO prediction for the FB asymmetry,

$$\Delta A_{FB}^X(m_{\ell\ell}) = A_{FB}^X(m_{\ell\ell}) - A_{FB}^{\text{LO}}(m_{\ell\ell}), \quad (6.12)$$

where $X = \text{NLO}_{\text{QCD}}$, NLO_{EW} or NNLO_{MIX} ⁶. In Fig. 6.5, we show results for $A_{FB}(m_{\ell\ell})$ with an increasing radiative content. The shifts compared to the LO asymmetry are shown in the bottom panels of Fig. 6.5.

We first observe that the NLO QCD corrections have a tiny impact throughout the whole invariant-mass range: the strong interaction is invariant under parity, and QCD corrections largely cancel in the definition of $A_{FB}(m_{\ell\ell})$. The NLO EW corrections have a large impact, through two distinct mechanisms: the large parity-even QED correction contributes mostly due to final-state radiation, which is particularly important below the Z resonance; the purely weak correction contributes instead to the parity-violating contribution. These two effects combined yield the largest shift of $A_{FB}(m_{\ell\ell})$. The mixed QCD–EW corrections inherit both features from the NLO EW terms, with increasing impact at larger $y_{\mu\mu}$ values; they are larger in size than the purely QCD corrections.

The effective weak mixing angle $\sin^2 \theta_{\text{eff}}^\ell$ can be determined from the study of $A_{FB}(m_{\ell\ell})$ [235]. In the Z -resonance region $A_{FB}(m_{\ell\ell})$ is generated by the product of vector and axial-vector couplings of

⁶Here NNLO_{MIX} includes only the mixed QCD–EW corrections on top of the LO prediction.




Figure 6.6: The variation of $A_{FB}(m_{\ell\ell})$ for a change of the weak mixing angle value by $19 \cdot 10^{-5}$.

the Z boson to fermions, in the squared Z -exchange diagrams, yielding in turn a sensitivity to the value of the weak mixing angle, present in the vector coupling. Away from the resonance, the squared Z -exchange diagrams are kinematically suppressed, $A_{FB}(m_{\ell\ell})$ is larger than at the resonance and driven by the product of the photon vector coupling with the axial coupling of the Z boson. This second coupling combination has no dependence on the weak mixing angle.

The $\sin^2 \theta_{\text{eff}}^\ell$ determination cannot be based on the analysis of one single bin of the $A_{FB}(m_{\ell\ell})$ invariant-mass distribution, i.e. the Z -resonance bin, because the result would be swamped by PDF uncertainties. The latter can be constrained [308], adopting a template fit procedure, which typically requires the analysis of $A_{FB}(m_{\ell\ell})$ in the mass window (70, 110) GeV. It is thus necessary to have control over the radiative corrections in this very region.

In order to illustrate the sensitivity of $A_{FB}(m_{\ell\ell})$ to $\sin^2 \theta_{\text{eff}}^\ell$, in Fig. 6.6 we study the variation of the asymmetry induced by a change⁷ of the weak mixing angle by $1.9 \cdot 10^{-4}$, a value slightly larger than the current experimental error on $\sin^2 \theta_{\text{eff}}^\ell$ [309]. For the sensitivity estimate, it is sufficient to consider a variation of the LO $A_{FB}(m_{\ell\ell})$ asymmetry, which we present in Fig. 6.6. The extraction of $\sin^2 \theta_{\text{eff}}^\ell$ from a fit to the experimental data must be performed using the most precise theoretical templates: missing higher orders could affect the final result inducing a spurious shift of the central value. The size of the latter can be estimated evaluating the impact of the higher order corrections on $A_{FB}(m_{\ell\ell})$.

Comparing Fig. 6.5 and Fig. 6.6, we clearly see that the impact of the mixed QCD–EW corrections is large. Part of the effect is known to be due to QED final-state radiation [214], which is typically accounted for in the experimental analyses using Monte Carlo parton showers. Nonetheless, when comparing Fig. 6.5 and Fig. 6.6, it is evident that a thorough evaluation of the computed corrections will be essential for accurate future determinations of the effective electroweak mixing angle $\sin^2 \theta_{\text{eff}}^\ell$.

⁷Since our calculation is performed with (G_μ, m_W, m_Z, m_H) as input parameters of the EW Lagrangian, we induce a variation of the weak mixing angle by a change of the value of the W -boson mass of 10 MeV.

6.3 Comparison with massless calculation

In this Section we present a comparison of our results to those of Ref. [217]. To this purpose, the EW scheme and input parameters are tuned to those of Ref. [217], as far as possible. In Ref. [217] the selection cuts are

$$m_{\ell\ell} > 200 \text{ GeV}, \quad p_{T,\ell^\pm} > 30 \text{ GeV}, \quad \sqrt{p_{T,\ell^+} p_{T,\ell^-}} > 35 \text{ GeV}, \quad |y_{\ell^\pm}| < 2.5, \quad (6.13)$$

and leptons are treated as massless.

It is well known that fiducial cuts applied to *bare* massless leptons are, in general, not collinear safe with respect to the emission of a collinear photon off a final-state lepton. The definition of physical observables requires therefore the introduction of a recombination procedure of photons to a close-by lepton⁸. This prevented the authors of Ref. [217] to directly compare their results, valid for *dressed* or *calo* leptons, to the results of Ref. [216], obtained for the case of bare massive muons.

Conversely, our framework, which retains the dependence on the mass of the final-state leptons, allows us to perform calculations for both bare and dressed leptons. This can be easily understood if we take seriously the fact that the mass of the lepton acts as the regulator of the final-state collinear singularity associated with the emission of a photon off a lepton. Indeed, the resulting dead-cone effect leads to a smooth suppression of the collinear photon radiation at the parton level. In combination with the $q_T/m_{\ell\ell} > r_{\text{cut}}$ requirement, applied at the parton/bare level, this makes the real-emission cross section finite. We can then apply the recombination procedure to each parton level event, analogously to what happens for jet cross sections computed with a local subtraction method. We also stress that the mass of the lepton should not be regarded as an additional cut-off parameter in our formalism (besides the slicing parameter r_{cut}). Our subtraction formula in Eq. (6.3) indeed retains the exact dependence on the mass of the lepton, without introducing further approximations.

To summarise, we can compare with the results of Ref. [217] for dressed leptons within our framework by repeating the calculation for decreasing lepton masses and eventually taking the limit $m_\ell \rightarrow 0$. In practice, we can perform the calculation at a sufficiently small value of the lepton mass, which can be even larger than the physical mass of the considered leptons, as long as power corrections can be safely neglected. Our comparison with the results for dressed electrons of Ref. [217] is carried out using a reference lepton mass $m_\ell = m_\mu = 105.658369 \text{ MeV}$.

The above procedure is consistent with the fact that we are performing an on-shell renormalization of the electric charge. Instead, one can consider a running coupling constant as, for example, in the $\overline{\text{MS}}$ scheme with only massless fermions (quarks and leptons) actively running in the beta function, while the wave functions of the external massive leptons are renormalized on-shell (decoupling scheme). In this situation, it is well known that an extra correction contribution, proportional to $\ln \mu_R/m_\ell$, must be added in passing from the massive to the massless calculation in order to compensate for the different number of active flavours in the beta function [310].

There is, however, a caveat related to LO photon-induced processes, namely the $\gamma\gamma \rightarrow \ell^+\ell^-$ channel. In this case, the subtraction of the photonic vacuum polarization insertions is usually dealt with relying on dimensional regularization for the light flavours, whereas, for massive fermions, the mass is used as the physical regulator. This is consistent with the reabsorption of initial-state collinear singularities into the renormalization of the parton distribution functions when applying collinear factorization, with the light flavours treated as active in the $\overline{\text{MS}}$ scheme, whereas the massive flavours are only generated radiatively.

⁸In Ref. [217] a lepton and a photon are recombined if their separation in rapidity and azimuth $R_{\ell\gamma} = \sqrt{\Delta^2 y_{\ell\gamma} + \Delta\varphi_{\ell\gamma}^2}$ is smaller than a parameter $R_{\text{cut}} = 0.1$

σ [pb]	σ_{LO}	$\sigma^{(1,0)}$	$\sigma^{(0,1)}$	$\sigma^{(2,0)}$	$\sigma^{(1,1)}$
$q\bar{q}$	1561.52(5)	340.3(3)	-49.77(5)	44.6(4)	-15.7(7)
qg	—	0.0601(3)	—	-32.7(2)	2.15(7)
$q\gamma$	—	—	-0.30(2)	—	-0.231(9)
$g\gamma$	—	—	—	—	0.266(1)
gg	—	—	—	2.02(6)	—
$\gamma\gamma$	59.645(6)	—	3.174(9)	—	—

Table 6.2: The different perturbative contributions to the fiducial cross section in the various partonic channels, in the setup of Ref. [217].

tively from scales above the pair production threshold. To perform a consistent comparison between our massive calculation and the massless results available in the literature, this difference in the collinear factorization must be compensated for. In other words, the photon distribution receives an $\mathcal{O}(\alpha)$ correction due to the change of scheme [311], and the cross section in the $\gamma\gamma$ channel is modified as

$$d\sigma_{\gamma\gamma, m_\ell=0}^{(0,1)} = d\sigma_{\gamma\gamma, m_\ell}^{(0,1)} - \alpha \frac{2e_\ell^2}{3\pi} \ln \frac{m_\ell^2}{\mu_F^2} d\sigma_{\gamma\gamma, m_\ell}^{(0,0)}, \quad (6.14)$$

which is obtained through the Abelianization of the analogous formula for the case of the gg channel given in Ref. [310]. From a practical point of view, the logarithmic term on the right-hand side represents the exact finite shift required to translate our massive computation into the $\overline{\text{MS}}$ factorization scheme adopted by the massless calculation. Subtracting this term from our massive result ensures a theoretically consistent, scheme-matched comparison between the two approaches.

In Tab. 6.2, we show our results for different orders in the different partonic channels. Our predictions for the NLO EW corrections are in very good agreement with those of Ref. [217]. The small difference for the quark–antiquark ($q\bar{q}$) channel⁹ is related to the fact we are using a finite lepton mass. Nonetheless, such a difference is well below the 0.01% level for fiducial cross sections, confirming that the power corrections in the lepton mass are indeed negligible for the chosen value.

Turning to higher orders, we observe a convincing agreement for the NNLO QCD corrections in all considered partonic channels with the corresponding results of Ref. [217]. The choice of *product cuts* [298] in Eq. (6.13) alleviates the problem of linear power corrections in the q_T subtraction formalism, leading to a better convergence of the method. Concerning the mixed QCD–EW corrections, we find agreement within uncertainties in the quark–antiquark channel, which contains the genuine two-loop virtual contribution, and in the gluon–photon channel. Our result for the correction in the quark–photon channel is about 10% larger than the corresponding result of Ref. [217]. We note that Ref. [217] quotes an overall 1% uncertainty on the mixed QCD–EW corrections, but does not provide the numerical accuracy of the separate partonic contributions, so we are unable to assess the significance of this small discrepancy. If the 10% difference in the quark–photon channel persists after acquiring higher Monte Carlo statistics, a more in-depth theoretical investigation will be necessary. Possible directions to look into would be differences in treatment of the renormalization scheme in the complex mass scheme or in the evaluation of the photon self-energy. In the quark–gluon channel, however, the discrepancy is definitely significant, our result being about a factor of two larger than the corresponding result of Ref. [217].

⁹For simplicity, the quark–antiquark channel is assumed to include all possible scattering processes involving quarkantiquark, quarkquark, or antiquarkantiquark initial states of any flavour.





Figure 6.7: Mixed QCD–EW corrections $\sigma^{(1,1)}$ as functions of the parameter r_{cut} in the quark–antiquark (top left), quark–gluon (bottom left), quark–photon (top right) and gluon–photon (bottom right) channels for two different values of the final-state lepton mass, $m_\ell = 1 \text{ GeV}$ and $m_\ell = m_\mu$. Our results are normalized to the $m_\ell = m_\mu$ result in the $r_{\text{cut}} \rightarrow 0$ limit.

Despite our efforts and two completely independent numerical implementations, we were unable to explain this difference. After extensive cross checks, the authors of Ref. [217] informed us that they found a bug in their implementation. After fixing this bug their result in the quark–gluon is in perfect agreement with ours. More detailed checks, also at the differential level, are needed for a thorough comparison of the two calculations.

In Fig. 6.7, we study the dependence on the slicing parameter r_{cut} for the nominal lepton mass $m_\ell = m_\mu$ and for a larger mass value, $m_\ell = 1 \text{ GeV}$. As expected, the cross-section exhibits a linear dependence on r_{cut} in the asymptotic limit. The results are consistent with the assumption that the residual dependence on the lepton mass, after the cancellation of large logarithms between real and virtual contributions, is indeed power suppressed. We observe that the power corrections of the lepton mass have a marginal impact on the final result. For comparison, the effect of increasing the lepton mass to $m_\ell = 1 \text{ GeV}$ on NLO EW corrections is to reduce them by about 3%. Our results provide a strong confirmation that the cancellation of all IR divergencies in our method does take place correctly, thereby providing a stringent check of the whole calculation.

We conclude this section by presenting further results for dressed leptons, which correspond to the setup of Sect. 6.2.1. In Table 6.3, we show the corresponding fiducial cross sections. Figure 6.8 displays the rapidity spectrum and the invariant mass distribution of the dilepton system. As discussed above, we use a fixed lepton mass, setting $m_l = m_\mu$, with negligible power corrections. As expected, the effect of QED final-state radiation is reduced when considering dressed leptons, leading to a decrease in radiative corrections in the NLO EW and mixed QCD-EW contributions. This effect is more pronounced in the invariant mass distribution, where the radiative tail below the resonance peak is significantly suppressed.

	σ [pb]	$\sigma^{(i,j)}$ [pb]	$\sigma^{(i,j)}/\sigma_{\text{LO}}$
LO	$735.73(3)^{+12.7\%}_{-13.6\%}$	—	—
NLO _{QCD}	$790.23(4)^{+2.7\%}_{-4.4\%}$	$\sigma^{(1,0)} = +54.50(5)$	+7.4%
NNLO _{QCD}	$774.4(7)^{+0.7\%}_{-0.7\%}$	$\sigma^{(2,0)} = -15.8(5)$	-2.2%
NLO _{EW}	$715.94(8)^{+12.7\%}_{-13.6\%}$	-19.7(1)	-2.7%
NNLO _{QCD} +NLO _{EW}	$754.6(5)^{+0.4\%}_{-0.6\%}$	—	—
NNLO _{QCD} +MIX+NLO _{EW}	$752.4(5)^{+0.7\%}_{-0.6\%}$	-2.1(1)	-0.3%

Table 6.3: As in Tab. 6.1 but for dressed leptons.**Figure 6.8:** As in Fig. 6.1 but for dressed leptons.

Constraining new physics with SMEFT global fits

“A blank spot on the map is a fine thing. It is a very fine thing indeed to see what is there, to go out into the unknown and discover.”

Eric Shipton, *English alpinist*.

In the preceding Chapters, we dedicated significant effort to pushing the boundaries of theoretical precision, culminating in the computation of fully differential cross-sections for the NC-DY process at NNLO QCD-EW level of accuracy. However, producing a high-precision SM prediction is not an end in itself, but rather an essential prerequisite for isolating the subtle and indirect signatures of BSM physics.

As introduced in Ch. 1, the Standard Model Effective Field Theory (SMEFT) provides a robust, model-independent framework to parameterize the effects of heavy, unknown new physics by adding to the SM Lagrangian higher-dimensional operators, scaled by unknown Wilson coefficients. Searching for new physics therefore becomes a parameter extraction exercise, i.e. we compute theoretical predictions as functions of these Wilson coefficients and perform a statistical fit to constrain them against precise experimental measurements. Crucially, the constraining power of such a global fit depends on the significance of the comparison, which combines both experimental and theoretical errors. If our predictions suffer from large theoretical uncertainties, such as scale uncertainties or missing higher-order corrections, the resulting error bands will hide the small deviations induced by BSM physics, rendering the fit insensitive to new phenomena. By reducing theoretical uncertainties through the multi-loop calculations detailed in this thesis, we ensure that theoretical errors do not become the bottleneck in this statistical comparison.

Furthermore, a single high-precision measurement is generally insufficient to bound the multidimensional parameter space of the SMEFT. To break the inherent degeneracies between various dimension-six operators, these highly precise theoretical predictions must be systematically confronted with a vast and diverse experimental dataset. Our goal in this final Chapter is to translate the theoretical machinery we have developed into statistically meaningful constraints on the fundamental laws of nature, extracting limits on Wilson coefficients using current data, and projecting how these bounds will evolve given the anticipated increase in experimental precision at future colliders. This Chapter has been adapted from [7].

7.1 The framework of SMEFT global fits

If the mass scale of new physics, Λ , is significantly larger than the electroweak scale ($v \approx 246$ GeV) and the typical kinematic reach of the collider, its effects can be systematically parametrized using an effective field theory approach. The SMEFT provides a model-independent parametrization of BSM effects by effectively extending the SM Lagrangian by means of higher-dimensional operators constructed

from SM fields that respect the $SU(3)_c \times SU(2)_L \times U(1)_Y$ local gauge symmetry. The extraction of meaningful physical constraints from experimental data requires a comprehensive and highly structured statistical approach known as a global fit.

The primary goal of a global fit is to simultaneously constrain a large set of SMEFT Wilson coefficients using a vast and diverse array of experimental measurements. Relying on single, isolated observables to constrain new physics is fundamentally limited by the presence of "flat directions" or parameter degeneracies. Because multiple dimension-6 operators can contribute to the same physical process, a deviation (or lack thereof) in a single measurement can often be explained by various combinations of Wilson coefficients canceling each other out.

To break these degeneracies, one must combine datasets with different kinematic dependencies and initial/final states. For instance, operators modifying the Z -boson couplings to fermions affect both pole electroweak precision observables at LEP and the high-energy tails of the Drell-Yan invariant mass distribution at the LHC. Only by performing a simultaneous global fit to both datasets can the individual contributions of these operators be disentangled and rigorously bounded.

7.1.1 SMEFIT code

In our study we used the publicly available fitting framework SMEFIT [312]. The latter takes as input some experimental datasets and their corresponding theoretical predictions. In particular, in the presence of n_{op} dimension-six SMEFT operators, a general SM cross-section σ_{SM} can be written as follows

$$\sigma_{\text{eff}}(\tilde{c}/\Lambda^2) = \sigma_{\text{SM}} + \sum_{i=1}^{n_{\text{op}}} \tilde{\sigma}_{\text{eff},i} \frac{\tilde{c}_i}{\Lambda^2} + \sum_{i,j=1}^{n_{\text{op}}} \tilde{\sigma}_{\text{eff},ij} \frac{\tilde{c}_i \tilde{c}_j}{\Lambda^4}, \quad (7.1)$$

with $\tilde{\sigma}_{\text{eff},i}/\Lambda^2$ and $\tilde{\sigma}_{\text{eff},ij}/\Lambda^4$ indicating respectively the contributions to the cross-section arising from the interference with the SM amplitudes and from the square of the EFT ones, once the Wilson coefficients $\tilde{c}_i$ are factored out. We explicitly note that the contributions originating from double insertions of dimension-six operators into the same amplitude, formally $\mathcal{O}(1/\Lambda^4)$, are not included in this parameterization. These cross-sections hence depend on n_{op} Wilson coefficients and on the cutoff scale Λ , with only the ratios $\tilde{c}/\Lambda^2$ being accessible in a model-independent analysis. For this reason, in the following we will reabsorb the scale Λ into the Wilson Coefficients by defining $c_i = \tilde{c}_i/\Lambda^2$. Eq. 7.1 can be generalized when other types of SMEFT operators, e.g. dimension-eight operators, are considered in the interpretation of the observable.

The terms σ_{SM} , $\tilde{\sigma}_{\text{eff},i}$ and $\tilde{\sigma}_{\text{eff},ij}$ in Eq. 7.1, are inputs to SMEFIT and are provided by means of external calculations. In this respect, the fitting code is agnostic in the calculational settings used to produce them, provided they comply with the required format of the theory tables.

The goal of SMEFIT is to determine confidence level in the space of the Wilson coefficients given an input dataset $\mathcal{D}$ and the corresponding theory predictions $\mathcal{T}(c)$, with the latter given by Eq. 7.1 or generalizations thereof. The agreement between the dataset $\mathcal{D}$ and a theory hypothesis $\mathcal{T}(c)$ is quantified by the likelihood function $\mathcal{L}(c)$. The wide majority of measurements used in EFT interpretations are presented as multi-Gaussian distributions, for which the likelihood is given by

$$\mathcal{L}(c) \propto \prod_{m,n=1}^{n_{\text{dat}}} \exp \left[- (\mathcal{T}_m(c) - \mathcal{D}_m) (\text{cov}^{-1})_{mn} (\mathcal{T}_n(c) - \mathcal{D}_n) \right]. \quad (7.2)$$

In such cases, the log-likelihood function $(-\log \mathcal{L})$ becomes either a quadratic or a quartic function of the Wilson coefficients $\mathbf{c}$, depending on whether the quadratic terms in the EFT parametrization of Eq. 7.1 are retained, according to the user specifications in the runcard.

Once a statistical model $\mathcal{L}(\mathbf{c})$ for the input dataset $\mathcal{D}$ is defined in terms of the theory predictions $\mathcal{T}(\mathbf{c})$, SMEFT proceeds to determine the most likely values of the Wilson coefficients $\mathbf{c}$ and the corresponding uncertainties by means of a strategy based on Nested Sampling (NS) via the MultiNest library [313,314]. The starting point is Bayes' theorem relating the posterior probability distribution of the parameters $\mathbf{c}$ given the observed data and the theory hypothesis, $P(\mathbf{c}|\mathcal{D}, \mathcal{T})$, to the likelihood function (conditional probability) and the prior distribution π ,

$$P(\mathbf{c}|\mathcal{D}, \mathcal{T}) = \frac{P(\mathcal{D}|\mathcal{T}, \mathbf{c}) P(\mathbf{c}|\mathcal{T})}{P(\mathcal{D}|\mathcal{T})} = \frac{\mathcal{L}(\mathbf{c}) \pi(\mathbf{c})}{\mathcal{Z}}, \quad (7.3)$$

with the Bayesian evidence $\mathcal{Z}$ ensuring the normalisation of the posterior distribution,

$$\mathcal{Z} = \int \mathcal{L}(\mathbf{c}) \pi(\mathbf{c}) d\mathbf{c}. \quad (7.4)$$

By means of Bayesian inference, NS maps the n_{op} -dimensional integral over the prior density into

$$X(\lambda) = \int_{\{\mathbf{c}: \mathcal{L}(\mathbf{c}) > \lambda\}} \pi(\mathbf{c}) d\mathbf{c}, \quad (7.5)$$

a one-dimensional function corresponding to the volume of the prior density $\pi(\mathbf{c})d\mathbf{c}$ associated to likelihood values greater than λ . Eq. 7.5 defines a transformation between the prior and posterior distributions sorted by the likelihood of each point in the EFT parameter space, is evaluated numerically, and results in n_{spl} samples $\{\mathbf{c}^{(k)}\}$ providing a representation of the posterior probability distribution from which one can evaluate confidence level intervals and related statistical estimators. The default SMEFT analysis assumes a flat prior volume $\pi(\mathbf{c})$, although implementing alternative functional forms for the prior volume is an option available to the user.

One benefit of sampling methods such as NS is that they bypass limitations of numerical optimisation techniques such as local minima preventing reaching the absolute minimum, with a drawback being that the computational resources required in NS grow exponentially with the dimensionality n_{op} .

The output of a SMEFT analysis consists on the posterior probability distributions associated to the fit coefficients $\mathbf{c}$ as well as ancillary statistical estimators such as the fit quality per dataset. This output can be processed and visualized by means of a fit report, which is generated by specifying a separate runcard either for an individual fit or for pair-wise (or multiple) comparisons between fits.

By modifying this dedicated runcard the user can specify what to display in the report, with currently available options including comparisons between SM and best-fit EFT predictions for individual datasets, posterior distributions with associated correlation and confidence level bar plots, two-parameter contour plots, the log-likelihood distribution among replicas or samples, Fisher matrix by dataset, and the outcome of the PCA analysis among others.

7.1.2 SMEFT Operator Basis

At mass dimension six, the SMEFT Lagrangian contains 2499 independent operators (including flavor and CP-violating structures, but not baryon and lepton number violating ones). In practice, it is impossible to constrain this entire parameter space simultaneously with current data. Therefore, a global fit must adopt a specific non-redundant basis and a well-defined truncation strategy.

In the following, we summarise the definitions of the SMEFT operator basis we used in our study. For each operator, we indicate its definition in terms of the SM fields, and the conventions that are used both for the operator and for the coefficient. For the predictions, we adopt the Warsaw basis of dimension-six operators [40] assuming a $U(2)_{q_L} \times U(2)_{u_R} \times U(3)_{d_R} \times (U(1)_\ell \times U(1)_e)^3$ flavour symmetry, subsequently restricted to $U(2)_{q_L} \times U(2)_{u_R} \times U(3)_{d_R} \times U(3)_\ell \times U(3)_e$ at the level of the global fits. This assumption corresponds to a $U(2)$ flavour symmetry framework [315,316]. This specific structure is motivated by the large top-quark mass, explicitly isolating the third-generation quarks from the first two light generations. In practice, this symmetry strongly restricts the structure of New Physics deviations in flavour-changing neutral currents since, in the exact symmetry limit, transitions between the light generations and the third generation are strictly forbidden.

- Table 7.1 lists the purely bosonic dimension-six SMEFT operators entering our basis. These operators modify the production and decay of Higgs bosons and the interactions of the gluons and electroweak gauge bosons. See [317] for more details.
- Table 7.2 lists the operators containing two fermion fields, either quarks or leptons. The two-fermion operators are classified in three groups: those involving 3rd generation quarks, those involving 1st and 2nd generation quarks, and those involving two leptons. For the latter (two lepton operators), the flavour index j runs from 1 to 3 following our flavour assumptions. Note that our basis includes the two-fermion operators modifying the Yukawa couplings of the top, bottom, and charm quarks ($c_{t\varphi}$, $c_{b\varphi}$, $c_{c\varphi}$) and of the tau and muon leptons ($c_{\tau\varphi}$, $c_{\mu\varphi}$). In Table 7.2, the coefficients indicated with (*) do not correspond to physical degrees of freedom in the fit, but are rather replaced by $c_{\varphi q}^{(-)}$, $c_{\varphi Q}^{(-)}$, and c_{tZ} .
- Table 7.3 collects the definition of the four-fermion coefficients that enter in the predictions in terms of the coefficients of Warsaw basis operators. These four fermion operators, involving either four-heavy quarks (4H), two-light-two-heavy quarks (2L2H), two-leptons-two-heavy quarks (2L2Q) or two-leptons-two-light quarks (2L2q), are defined as

$$\begin{aligned}
\mathcal{O}_{qq}^{1(ijkl)} &= (\bar{q}_i \gamma^\mu q_j) (\bar{q}_k \gamma_\mu q_l), \\
\mathcal{O}_{qq}^{3(ijkl)} &= (\bar{q}_i \gamma^\mu \tau^I q_j) (\bar{q}_k \gamma_\mu \tau^I q_l), \\
\mathcal{O}_{qu}^{1(ijkl)} &= (\bar{q}_i \gamma^\mu q_j) (\bar{u}_k \gamma_\mu u_l), \\
\mathcal{O}_{qu}^{8(ijkl)} &= (\bar{q}_i \gamma^\mu T^A q_j) (\bar{u}_k \gamma_\mu T^A u_l), \\
\mathcal{O}_{qd}^{1(ijkl)} &= (\bar{q}_i \gamma^\mu q_j) (\bar{d}_k \gamma_\mu d_l), \\
\mathcal{O}_{qd}^{8(ijkl)} &= (\bar{q}_i \gamma^\mu T^A q_j) (\bar{d}_k \gamma_\mu T^A d_l), \\
\mathcal{O}_{uu}^{(ijkl)} &= (\bar{u}_i \gamma^\mu u_j) (\bar{u}_k \gamma_\mu u_l), \\
\mathcal{O}_{ud}^{1(ijkl)} &= (\bar{u}_i \gamma^\mu u_j) (\bar{d}_k \gamma_\mu d_l), \\
\mathcal{O}_{ud}^{8(ijkl)} &= (\bar{u}_i \gamma^\mu T^A u_j) (\bar{d}_k \gamma_\mu T^A d_l).
\end{aligned} \tag{7.6}$$

Moreover we include the four-lepton (4l) operators

$$\begin{aligned}
\mathcal{O}_{\ell\ell}^{(jklm)} &= (\bar{\ell}_j \gamma^\mu \ell_k) (\bar{\ell}_l \gamma^\mu \ell_m), \\
\mathcal{O}_{\ell e}^{(jklm)} &= (\bar{\ell}_j \gamma^\mu \ell_k) (\bar{e}_l \gamma^\mu e_m), \\
\mathcal{O}_{ee}^{(jklm)} &= (\bar{e}_j \gamma^\mu e_k) (\bar{e}_l \gamma^\mu e_m).
\end{aligned} \tag{7.7}$$

Operator	Coefficient	Definition	Operator	Coefficient	Definition
$\mathcal{O}_{\varphi G}$	$c_{\varphi G}$	$(\varphi^\dagger \varphi) G_A^{\mu\nu} G_{\mu\nu}^A$	$\mathcal{O}_{\varphi \square}$	$c_{\varphi \square}$	$(\varphi^\dagger \varphi) \square (\varphi^\dagger \varphi)$
$\mathcal{O}_{\varphi B}$	$c_{\varphi B}$	$(\varphi^\dagger \varphi) B^{\mu\nu} B_{\mu\nu}$	$\mathcal{O}_{\varphi D}$	$c_{\varphi D}$	$(\varphi^\dagger D^\mu \varphi)^\dagger (\varphi^\dagger D_\mu \varphi)$
$\mathcal{O}_{\varphi W}$	$c_{\varphi W}$	$(\varphi^\dagger \varphi) W_I^{\mu\nu} W_{\mu\nu}^I$	$\mathcal{O}_W$	c_{WWW}	$\epsilon_{IJK} W_{\mu\nu}^I W^{J,\nu\rho} W_\rho^{K,\mu}$
$\mathcal{O}_{\varphi WB}$	$c_{\varphi WB}$	$(\varphi^\dagger \tau_I \varphi) B^{\mu\nu} W_{\mu\nu}^I$	$\mathcal{O}_\varphi$	c_φ	$(\varphi^\dagger \varphi)^3$

Table 7.1: The purely bosonic dimension-six SMEFT operators entering our basis. These operators modify the production and decay of Higgs bosons and the interactions of the electroweak gauge bosons. For each operator, we indicate its definition in terms of the SM fields, and the conventions that are used both for the operator and for the coefficient. See [317] for more details.

The flavour index i is either 1 or 2, and j is either 1, 2 or 3.

In the global fits which will be presented in Sect. 7.2 we impose a $U(3)_\ell \times U(3)_e$ symmetry in the lepton sector. In total, the number of independent degrees of freedom is $n_{\text{op}} = 61$. The flavour universal coefficients definitions are trivially derived by identifying all lepton flavour indices. The four-lepton operator $\mathcal{O}_{\ell\ell}^{(jklm)}$, is split into two invariants, such that [318]

$$c_{\ell\ell}^{jklm} = c_{\ell\ell} \delta^{jk} \delta^{lm} + c'_{\ell\ell} \delta^{jm} \delta^{lk}. \quad (7.8)$$

The coefficients used in the fit are classified into the different groups of operators as specified in Table 7.4. For each operator class, we indicate the number of members and the degrees of freedom considered in the SMEFT global fit, see Tables 7.1–7.3 for the corresponding definitions.

7.1.3 Experimental datasets and Drell-Yan

The LHC dataset used in this work follows the one adopted in [319, 320] with two extensions. First, the inclusion of Drell-Yan production, specifically of the CMS 13 TeV measurement of [301] which is particularly relevant to constrain the four-fermion operators composed of two quark and two leptonic fields, namely:

$$c_{Q\ell}^{(-)}, c_{Q\ell}^{(3)}, c_{Qe}, c_{t\ell}, c_{te}, c_{q\ell}^{(-)}, c_{q\ell}^{(3)}, c_{qe}, c_{lu}, c_{ld}, c_{eu}, c_{ed}. \quad (7.9)$$

In particular, $c_{Q\ell}^{(-)}, c_{Q\ell}^{(3)}, c_{Qe}$, enter at LO through the $b\bar{b}$ channel, while $c_{t\ell}$ and c_{te} start contributing at NLO. Second, the ATLAS measurement of Z production in vector boson fusion (VBF) at 13 TeV [321], which provides independent constraints on the modifications of the Triple Gauge Couplings.

The CMS Z' dilepton search of [301] can be recasted as a measurement of neutral-current DY production. This analysis targets the high invariant mass region for the dilepton pair, which is particularly sensitive to two-lepton-two-quark operators due to energy growing effects [322, 323]. The CMS dataset, in particular, provides the full Run-2 statistics and extends the dilepton invariant-mass reach up to 6 TeV [324]. In our study, we retain the $n_b = 79$ bins with $m_{\ell\ell} \geq 500$ GeV. For each bin i , the CMS analysis reports the observed dilepton yield, $n_{i,\ell\ell}$, as well as the corresponding SM expectation $b_{i,\text{SM}}$, which contains contributions from DY, top, and diboson production. To isolate the DY signal, n_i , we

Operator	Coefficient	Definition	Operator	Coefficient	Definition
3rd generation quarks					
$\mathcal{O}_{\varphi Q}^{(1)}$	$c_{\varphi Q}^{(1)} (*)$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{Q} \gamma^\mu Q)$	$\mathcal{O}_{tW}$	c_{tW}	$i(\bar{Q} \tau^{\mu\nu} \tau_I t) \tilde{\varphi} W_{\mu\nu}^I + \text{h.c.}$
$\mathcal{O}_{\varphi Q}^{(3)}$	$c_{\varphi Q}^{(3)}$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \tau_I \varphi)(\bar{Q} \gamma^\mu \tau^I Q)$	$\mathcal{O}_{tB}$	$c_{tB} (*)$	$i(\bar{Q} \tau^{\mu\nu} t) \tilde{\varphi} B_{\mu\nu} + \text{h.c.}$
	$c_{\varphi Q}^{(-)}$	$c_{\varphi Q}^{(1)} - c_{\varphi Q}^{(3)}$		c_{tZ}	$-s_\theta c_{tB} + c_\theta c_{tW}$
$\mathcal{O}_{\varphi t}$	$c_{\varphi t}$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{t} \gamma^\mu t)$	$\mathcal{O}_{tG}$	c_{tG}	$igs(\bar{Q} \tau^{\mu\nu} T_A t) \tilde{\varphi} G_{\mu\nu}^A + \text{h.c.}$
$\mathcal{O}_{t\varphi}$	$c_{t\varphi}$	$(\varphi^\dagger \varphi) \bar{Q} t \tilde{\varphi} + \text{h.c.}$	$\mathcal{O}_{b\varphi}$	$c_{b\varphi}$	$(\varphi^\dagger \varphi) \bar{Q} b \varphi + \text{h.c.}$
1st, 2nd generation quarks					
$\mathcal{O}_{\varphi q}^{(1)}$	$c_{\varphi q}^{(1)} (*)$	$\sum_{i=1,2} i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{q}_i \gamma^\mu q_i)$	$\mathcal{O}_{\varphi u}$	$c_{\varphi u}$	$\sum_{i=1,2} i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{u}_i \gamma^\mu u_i)$
$\mathcal{O}_{\varphi q}^{(3)}$	$c_{\varphi q}^{(3)}$	$\sum_{i=1,2} i(\varphi^\dagger \overleftrightarrow{D}_\mu \tau_I \varphi)(\bar{q}_i \gamma^\mu \tau^I q_i)$	$\mathcal{O}_{\varphi d}$	$c_{\varphi d}$	$\sum_{j=1,2,3} i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{d}_j \gamma^\mu d_j)$
	$c_{\varphi q}^{(-)}$	$c_{\varphi q}^{(1)} - c_{\varphi q}^{(3)}$	$\mathcal{O}_{c\varphi}$	$c_{c\varphi}$	$(\varphi^\dagger \varphi) \bar{q}_2 c \tilde{\varphi} + \text{h.c.}$
two-leptons					
$\mathcal{O}_{\varphi \ell_j}$	$c_{\varphi \ell_j}$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{\ell}_j \gamma^\mu \ell_j)$	$\mathcal{O}_{\varphi e}$	$c_{\varphi e}$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{e} \gamma^\mu e)$
$\mathcal{O}_{\varphi \ell_j}^{(3)}$	$c_{\varphi \ell_j}^{(3)}$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \tau_I \varphi)(\bar{\ell}_j \gamma^\mu \tau^I \ell_j)$	$\mathcal{O}_{\varphi \mu}$	$c_{\varphi \mu}$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{\mu} \gamma^\mu \mu)$
$\mathcal{O}_{\mu\varphi}$	$c_{\mu\varphi}$	$(\varphi^\dagger \varphi) \bar{\ell}_2 \mu \varphi + \text{h.c.}$	$\mathcal{O}_{\varphi\tau}$	$c_{\varphi\tau}$	$i(\varphi^\dagger \overleftrightarrow{D}_\mu \varphi)(\bar{\tau} \gamma^\mu \tau)$
$\mathcal{O}_{\tau\varphi}$	$c_{\tau\varphi}$	$(\varphi^\dagger \varphi) \bar{\ell}_3 \tau \varphi + \text{h.c.}$			

Table 7.2: Same as Table 7.1 for the operators containing two fermion fields, either quarks or leptons. They are classified in three groups: those involving 3rd generation quarks, those involving 1st and 2nd generation quarks, and those involving two leptons. For the latter (two-lepton operators), the flavour index j runs from 1 to 3: with the flavour assumptions adopted in this work, these coefficients will be the same regardless of the specific values that j takes. The coefficients indicated with (*) do not correspond to degrees of freedom in the fit, but are rather replaced by $c_{\varphi q}^{(-)}$, $c_{\varphi Q}^{(-)}$, and c_{tZ} .

DoF	Definition (in Warsaw basis notation)	DoF	Definition (in Warsaw basis notation)
4-heavy-quark operators			
c_{QQ}^1	$2c_{qq}^{1(3333)} - \frac{2}{3}c_{qq}^{3(3333)}$	c_{QQ}^8	$8c_{qq}^{3(3333)}$
c_{Qt}^1	$c_{qu}^{1(3333)}$	c_{Qt}^8	$c_{qu}^{8(3333)}$
c_{tt}^1	$c_{uu}^{1(3333)}$		
2-light-2-heavy quark operators			
$c_{Qq}^{1,8}$	$c_{qq}^{1(ii3i)} + 3c_{qq}^{3(ii3i)}$	$c_{Qq}^{1,1}$	$c_{qq}^{1(ii33)} + \frac{1}{6}c_{qq}^{1(ii3i)} + \frac{1}{2}c_{qq}^{3(ii3i)}$
$c_{Qq}^{3,8}$	$c_{qq}^{1(ii3i)} - c_{qq}^{3(ii3i)}$	$c_{Qq}^{3,1}$	$c_{qq}^{3(ii33)} + \frac{1}{6}(c_{qq}^{1(ii33)} - c_{qq}^{3(ii3i)})$
c_{tq}^8	$c_{qu}^{8(ii33)}$	c_{tq}^1	$c_{qu}^{1(ii33)}$
c_{tu}^8	$2c_{uu}^{(ii33)}$	c_{tu}^1	$c_{uu}^{(ii33)} + \frac{1}{3}c_{uu}^{(i33i)}$
c_{Qu}^8	$c_{qu}^{8(33ii)}$	c_{Qu}^1	$c_{qu}^{1(33ii)}$
c_{td}^8	$c_{ud}^{8(33jj)}$	c_{td}^1	$c_{ud}^{1(33jj)}$
c_{Qd}^8	$c_{qd}^{8(33jj)}$	c_{Qd}^1	$c_{qd}^{1(33jj)}$
2-heavy-quark-2-lepton operators			
$c_{Q\ell_1}^{(-)}$	$c_{\ell q}^{1(1133)} - c_{\ell q}^{3(1133)}$	$c_{Q\ell_1}^{(3)}$	$c_{\ell q}^{3(1133)}$
$c_{Q\ell_3}^{(-)}$	$c_{\ell q}^{1(3333)} - c_{\ell q}^{3(3333)}$	$c_{Q\ell_3}^{(3)}$	$c_{\ell q}^{3(3333)}$
c_{Qe}	$c_{qe}^{(3311)}$	$c_{Q\tau}$	$c_{qe}^{(3333)}$
c_{te}	$c_{eu}^{(1133)}$	$c_{t\tau}$	$c_{eu}^{(3333)}$
$c_{t\ell_j}$	$c_{\ell u}^{(jj33)}$		
2-light-quark-2-lepton operators			
$c_{q\ell_1}^{(-)}$	$c_{\ell q}^{1(11ii)} - c_{\ell q}^{3(11ii)}$	$c_{q\ell_1}^{(3)}$	$c_{\ell q}^{3(11ii)}$
c_{qe}	$c_{qe}^{(ii11)}$		
$c_{\ell_1 u}$	$c_{\ell u}^{(11ii)}$	$c_{\ell_1 d}$	$c_{\ell d}^{(11jj)}$
c_{eu}	$c_{eu}^{(11ii)}$	c_{ed}	$c_{ed}^{(11jj)}$

Table 7.3: Definition of the four-fermion coefficients that enter in the fit in terms of the coefficients of Warsaw basis operators of Eq. (7.7). These coefficients are classified into four-heavy-quark, and two-light-two-heavy-quark, 2-heavy-quark-2-lepton, and 2-light-quark-2-lepton operators. The flavour index i is either 1 or 2, and j is either 1, 2 or 3: with the flavour assumptions adopted in this work, these coefficients will be the same regardless of the specific values that i and j take.

Operator class	Subclass	n_{op}	Wilson Coefficients
Purely bosonic (B)		8	$c_{\varphi G}, c_{\varphi B}, c_{\varphi \square}, c_{\varphi D}, c_{WWW}, c_{\varphi}, c_{\varphi W}, c_{\varphi WB}$
Two-fermion (2F)	3rd gen quarks	8	$c_{\varphi Q}^{(1)}, c_{tW}, c_{\varphi Q}^{(3)}, c_{tB}, c_{\varphi t}, c_{tG}, c_{t\varphi}, c_{b\varphi}$
	1st, 2nd gen quarks	5	$c_{\varphi q}^{(1)}, c_{\varphi q}^{(3)}, c_{\varphi d}, c_{\varphi u}, c_{c\varphi}$
	two-lepton (2ℓ)	5	$c_{\varphi \ell}, c_{\varphi \ell}^{(3)}, c_{\varphi e}, c_{\mu\varphi}, c_{\tau\varphi}$
Four-fermion (4F)	4H	5	$c_{QQ}^1, c_{QQ}^8, c_{Qt}^1, c_{Qt}^8, c_{tt}^1$
	2L2H	14	$c_{Qq}^{1,8}, c_{Qq}^{1,1}, c_{Qq}^{3,8}, c_{Qq}^{3,1}, c_{tq}^8, c_{tq}^1, c_{tu}^8, c_{tu}^1, c_{Qu}^8, c_{Qu}^1, c_{td}^8, c_{td}^1, c_{Qd}^8, c_{Qd}^1$
	2H2 ℓ	5	$c_{Q\ell}^{(-)}, c_{Q\ell}^{(3)}, c_{Qe}, c_{t\ell}, c_{te}$
	2L2 ℓ	7	$c_{q\ell}^{(-)}, c_{q\ell}^{(3)}, c_{qe}, c_{lu}, c_{ld}, c_{eu}, c_{ed}$
	4 ℓ	4	$c_{\ell\ell}, c'_{\ell\ell}, c_{\ell e}, c_{ee}$
Total		61	—

Table 7.4: Overview of Wilson coefficients associated to the dimension-six SMEFT operators that compose our fitting basis. For each operator class, we indicate the number of members and the degrees of freedom considered in the SMEFT global fit. See Tables 7.1–7.3 for the corresponding definitions. Here ‘H’ stands for a heavy quark field, L for a light quark field, and ℓ for a leptonic field. The index i for the 2-lepton operators runs from 1 to 3 in our flavour assumptions.

subtract the non-DY background (top and diboson) from the observed yield,

$$n_i = n_{i,\ell\ell} - b_i, \quad (7.10)$$

where the non-DY background, b_i , is obtained as

$$b_i = b_{i,\text{SM}} - a_i d_{i,\text{SM}}. \quad (7.11)$$

Here $d_{i,\text{SM}}$ denotes our prediction for the DY event yields based on MADGRAPH5_AMC@NLO [20], and a_i represents the experimental acceptance and selection efficiency. Theoretical predictions for the DY differential distributions in the SMEFT, $d_{i,\text{SMEFT}}$, have been computed using MADGRAPH5_AMC@NLO together with SMEFTSIM [318] and SMEFTATNLO [325], the latter by means of a dedicated version extended to two-lepton-two-quark (2L2Q) operators, for the LO and NLO QCD calculations respectively. Several cross-checks have been performed between SMEFTSIM and our new implementation of 2L2Q operators in SMEFTATNLO, finding agreement. This way, we obtain the expected number of events in the i -th bin as a function of the Wilson coefficients as $y_i = a_i d_{i,\text{SMEFT}}$.

Since we focus on the tails of the distributions, where a small number of events is observed, for this dataset the likelihood is non-Gaussian and is instead given by the product of Poisson probabilities per bin:

$$\mathcal{L}(n_i, y_i) = \prod_{i=1}^m \frac{y_i^{n_i}}{n_i!} e^{-y_i}. \quad (7.12)$$

The best-fit values for the Wilson coefficients are then determined by minimising the difference between the experimental and theoretical likelihoods,

$$D = -2 (\log \mathcal{L}(n_i, y_i) - \log \mathcal{L}(n_i, n_i)) = 2 \sum_{i=1}^m \left(y_i - n_i + n_i \log \frac{n_i}{y_i} \right), \quad (7.13)$$

through the external likelihood module of SMEFIT. We note that the Poissonian likelihood of Eq. 7.12 introduces non-linear effects even in the $\mathcal{O}(\Lambda^{-2})$ fits, and hence one may find non-Gaussian posteriors in linear EFT fits which include this DY dataset. In practice, we find that deviations from Gaussianity are small.

Concerning the Zjj production in vector boson fusion from ATLAS [321] at 13 TeV, here we include the $n_{\text{dat}} = 12$ bins of the absolute differential distribution in $\Delta\varphi_{jj}$, namely the difference in azimuthal angle between the two tagged VBF jets of the event. This distribution is known to provide good sensitivity to modifications in the triple-gauge coupling c_{WWW} [326]. For this dataset, theoretical SMEFT predictions have been computed at LO in MADGRAPH5_AMC@NLO interfaced with SMEFTATNLO.

7.1.4 Projections for lepton colliders

In our study, we consider also projections for future lepton collider observables. The full dataset, including theoretical SM and SMEFT predictions, and the projected experimental and theoretical uncertainties, are publicly available in the SMEFIT database, https://github.com/LHCfitNikhef/smefit_database, to which we refer the reader for more detailed information.

Table 7.5 summarizes the running scenarios assumed for the projections considered in this work following the European Strategy for Particle Physics Update 2026 (ESPPU26) submitted inputs. For each collider, we indicate the geometry, the number of interaction points (IPs), the beam polarisation, the centre of mass energy, and the integrated luminosity.

Collider	Geometry, # IPs	Beam Polarisation (e^-, e^+)	Energy ($\sqrt{s}$)	Luminosity ($\mathcal{L}_{\text{int}}$)
LEP3	Circular, 2 IPs	Unpolarised	91.2 GeV	48 ab^{-1}
			160 GeV	5.6 ab^{-1}
			230 GeV	2.304 ab^{-1}
FCC-ee	Circular, 4 IPs	Unpolarised	91.2 GeV	205 ab^{-1}
			161 GeV	19.2 ab^{-1}
			240 GeV	10.8 ab^{-1}
			365 GeV	3.12 ab^{-1}
LCF	Linear, 1 IP	$(\mp 80\%, \pm 30\%)$	91.2 GeV	0.1 ab^{-1}
			250 GeV	3 ab^{-1}
			350 GeV	0.2 ab^{-1}
		$(\mp 80\%, \pm 20\%)$	550 GeV	4 (8) ab^{-1}
			1000 GeV	8 (0) ab^{-1}

Table 7.5: The assumed running scenarios for the LEP3, FCC-ee and LCF colliders. In the case of LCF, the luminosities correspond to the LCF550 program, while those in parentheses correspond to LCF1000 (see text for the details).

LEP3. LEP3 is a proposed electron-positron collider operating as an electroweak and Higgs factory in the LHC tunnel. Here we follow [327] and assume three different runs at three different $\sqrt{s}$ and two IPs for the operation of LEP3. The Z -pole run at 91.2 GeV would run for 5 years at an integrated luminosity per year and per IP of 4.8 ab^{-1} , resulting in a total luminosity of 48 ab^{-1} . The WW threshold run at 160 GeV would last 4 years, collecting a total of 5.6 ab^{-1} . The Higgs factory ZH run at 230 GeV would collect a total of 2.304 ab^{-1} . The statistical uncertainties for LEP3 observables are obtained by rescaling the corresponding ones from the FCC-ee by the luminosity ratio $\sqrt{\mathcal{L}_{\text{LEP3}}/\mathcal{L}_{\text{FCC-ee}}}$, while the systematic uncertainties are kept the same as for FCC-ee (though we note that dedicated detector simulations are missing for LEP3). In terms of input observables, for LEP3 we consider the same processes as for FCC-ee, with the exception of those belonging to the $\sqrt{s} = 365$ GeV run which are not accessible at LEP3.

FCC-ee. In the integrated FCC program [328,329], an electron-positron collider (FCC-ee) operating in a new 91 km tunnel would be followed by an hadron-hadron collider (FCC-hh) with centre of mass energies between 70 TeV and 120 TeV, depending on the chosen magnet technology. The FCC-ee inputs used in this analysis are described in [330] and have been updated following the Feasibility Report [328, 329] studies submitted to the ESPPU26, with projected experimental uncertainties taken from [331,332].

The Z -pole measurements considered for the FCC-ee are the following: the electroweak coupling constant $\alpha(m_Z)$ and the Electroweak Precision Observables (EWPOs), namely the total Z width Γ_Z , the total cross section into hadrons σ_{had}^0 , the ratios to difermion final states $R_e, R_\mu, R_\tau, R_b, R_c$, and the corresponding forward-backward asymmetries $A_e, A_\mu, A_\tau, A_b, A_c$.

For the data-taking at $\sqrt{s} = 240$ GeV and $\sqrt{s} = 365$ GeV, we consider projections for single Higgs production, difermion production, and the total cross section into hadrons σ_{had}^0 . Specifically, for single Higgs production, we consider the HZ channel (both inclusive cross section and its decays to the $b\bar{b}$, $c\bar{c}$, gg , WW , ZZ , γZ , $\gamma\gamma \tau^+\tau^-$ and $\mu^+\mu^-$ final states) and the VBF channel (only hadronic Higgs decays at 240 GeV, and all channels except $\mu^+\mu^-$ at 365 GeV). For difermion production, we include both the forward-backward asymmetries ($A_e, A_\mu, A_\tau, A_b, A_c$) and the associated partial ratios (R_μ, R_τ, R_b, R_c), except for the e^+e^- final state which is included as total production rate (in addition to A_e),

Additional constraints provided by the FCC-ee projections arise from W^+W^- production (total rate, branching fractions, optimal observables for the fully- and semi-leptonic final states) at 161 GeV, 240 GeV, and 365 GeV. We also consider the optimal observables for $t\bar{t}$ production at $\sqrt{s} = 365$ GeV. For the latter, a 10% selection efficiency to reconstruct the $t\bar{t}$ system is assumed.

LCF. Several possible realisations of a linear electron-positron collider, from the International Linear Collider (ILC) in Japan to the Compact Linear Collider (CLIC) at CERN have been proposed for consideration. Here we adopt as a benchmark the settings of the Linear Collider Facility (LCF) vision for CERN's future [333]. As indicated in Table 7.5, the LCF would run at five different energies, from a Giga- Z run at 91.2 GeV and Higgs and top quark factory runs at 250 GeV and 350 GeV to the high-energy runs at 550 GeV and 1 TeV. An alternative running scenario is also being considered, where the 1 TeV run is replaced by a longer 550 GeV run collecting 8 ab^{-1} of luminosity. In this work, these scenarios are labelled as LCF550 and LCF1000. In Table 7.5 we indicate in brackets the luminosities corresponding to LCF550.

A unique feature of linear colliders is the possibility of polarising the lepton beams, which provides additional sensitivity to the electroweak couplings. The polarised e^- and e^+ beams of the LCF with polarisations $P = (P_-, P_+)$ contain the following fractions of right and left-handed particles [334]:

$$f_R(e^\pm) = \frac{1 + P_\pm}{2}, \quad f_L(e^\pm) = \frac{1 - P_\pm}{2}, \quad (7.14)$$

e.g. a beam with polarisation $P = (+80\%, -20\%)$ results in handedness fractions of

$$f_R(e^-) = 90\%, \quad f_L(e^-) = 10\%, \quad f_R(e^+) = 40\%, \quad f_L(e^+) = 60\%, \quad (7.15)$$

and hence collisions dominated by the $e_R^- + e_L^+$ initial state. Since right- and left-handed fermions are charged differently under the $SU(2)_L \otimes U(1)_Y$ gauge group, varying polarisations probe different aspects of the electroweak interactions, enhancing sensitivity to BSM deviations in electroweak couplings.

In terms of experimental input, for the LCF projections considered in this work we include the following observables:

- $\alpha(m_Z)$ and EWPOs, measured at $\sqrt{s} = 250$ GeV via radiative return. EWPOs are also measured in the Giga- Z run at 91 GeV.
- Single Higgs production in the ZH channel, measured both inclusively and in different decay channels at $\sqrt{s} = 250, 350$ and 550 GeV.
- Single Higgs production in the VBF channel, accessible in the different final states at 350, 550 and 1000 GeV. For the run at $\sqrt{s} = 250$ GeV, only the $b\bar{b}$ final state is available.
- The total rate and forward-backward asymmetries in difermion production, measured at all energies.

- W^+W^- optimal observables, included for all energies.
- $t\bar{t}$ optimal observables, included for all runs with $\sqrt{s} \geq 350$ GeV, with a 10% selection efficiency assumed up to 550 GeV and then a 6% efficiency for the 1000 GeV run.
- The high-energy runs at $\sqrt{s} = 550$ GeV and 1 TeV give access to $t\bar{t}H$ and double Higgs production, providing direct sensitivity to the top Yukawa coupling and to the Higgs self-coupling, respectively. Here we include the total inclusive cross sections for $t\bar{t}H$ and HH production at these two energies, for the latter considering both VBF (W^+W^- fusion) and associated ZHH production.

7.1.5 Theoretical uncertainties

In previous SMEFT analyses [319, 330] they considered two scenarios for the theoretical uncertainties affecting projections for EWPOs at future lepton colliders: the current level of theory errors, and no theory errors at all. Here we expand this treatment to take into account two additional scenarios for theory uncertainties in SM predictions discussed in the Physics Briefing Book (PBB) [335], namely the so-called ‘conservative’ and ‘aggressive’ scenarios. In the following, we consider SM theory uncertainties for the following lepton collider observables: Higgs production and decay, Z pole EWPOs, WW production, and W branching ratios. The four scenarios for SM theory uncertainties for lepton colliders that are used for this work are:

- **Current:** theoretical uncertainties on SM processes at future lepton colliders are assumed to be of the same size as current estimates.
- **Conservative:** theoretical uncertainties on lepton collider processes are assumed to improve by a moderate amount, based on the expansion of existing higher-order computational techniques.
- **Aggressive:** we assume a substantial improvement of theoretical uncertainties, which would require breakthroughs in higher-order calculation methods.
- **Ideal:** we assume all SM theory uncertainties at future colliders are sub-dominant compared to experimental precision and can be neglected.

In addition to theoretical uncertainties from the PBB, we include for Z pole observables (EWPOs and α) the effect of parametric uncertainties arising from the propagation of experimental uncertainty on the measurement of input parameters. The impact of parametric uncertainties is highly correlated, but these correlations were neglected in previous SMEFT analyses. Here we take them into account by including their full theory covariance. This was done using the code presented in [336] for LEP and for future colliders, by implementing the projected uncertainties on input parameters. In the G_μ scheme, these are $\{m_W, m_Z, G_\mu, m_H, m_t, \alpha_S\}$.

7.2 Global SMEFT analysis

In this section we present results for the global SMEFT fit at the three e^+e^- colliders listed in Table 7.5: LEP3, FCC-ee, and LCF. For LCF, we consider both the LCF550 (without the 1 TeV run, but with higher luminosity at $\sqrt{s} = 550$ GeV) and the LCF1000 variants separately. While comparing the projections from these three colliders among them, we also quantify the impact of quadratic EFT corrections, how the

sensitivity to SMEFT operators varies in the different scenarios on the theoretical uncertainties associated to the SM predictions, and revisit the study of [319] for the Higgs self-coupling at the HL-LHC and future electron-positron colliders.

In this section, we show results for the $n_{\text{op}} = 61$ Wilson coefficients that compose our fitting basis evaluated at the reference scale $\mu_0 = 10$ TeV. The reference scale μ_0 is introduced by renormalization procedure. While physical observables do not depend on it, the Wilson coefficients do, and their scale dependence is determined by Renormalization Group Equations (RGEs). We set $\mu_0 = 10$ TeV to reflect a possible mass scale of heavy new physics, without committing to any particular UV model. This minimizes large logarithms when matching our EFT results onto explicit UV-complete models. Choosing a lower scale motivated by high-precision measurements, such as m_Z would leave the physical observables invariant but shift the mathematical representation of the fit. Applying the RGE evolution matrix to translate our results down to m_Z would rotate the basis: central values would shift, errors would redistribute, and our single-parameter individual fits would appear as highly correlated, multi-operator scenarios at the lower scale. Furthermore, we present our results both in global fits, where all n_{op} Wilson coefficients are allowed to float simultaneously and are subsequently marginalised, as well as in individual fits. In the latter, only a single Wilson coefficient is assumed to be non-zero at the high scale $\mu_0 = 10$ TeV, though the coefficients of additional operators may be induced at the experimental scale via RGE mixing.

First, we present in Sect. 7.2.1 results for the baseline fits, carried out at the linear level in the EFT expansion and account for theoretical uncertainties in the aggressive scenario. Then we study in Sect. 7.2.2 the impact of quadratic EFT corrections in the global fit. Subsequently, Sect. 7.2.3 presents results for the four scenarios for the theoretical uncertainties discussed in Sect. 7.1.5.

7.2.1 Baseline results

For the presentation of the results in this section, we cluster the Wilson coefficients into different families: those associated with four-heavy-quark (4H), two-light-two-heavy quark (2L2H), purely bosonic (B), two-lepton-two-light-quark (2 ℓ 2q), two-fermion (2FB), two-lepton-two-heavy-quark (2 ℓ 2Q), and four-lepton (4 ℓ) operators. We start by presenting our baseline results for the global SMEFT fit, which are obtained with settings aligned whenever possible with those of the ESPPU26 Physics Briefing Book [335]. In particular, here we consider as baseline linear EFT calculations and the aggressive scenario for theoretical uncertainties at lepton colliders, and subsequently explore the impact of varying these settings.

To begin with, Fig. 7.1 displays the 95% C.I. bounds obtained for the Wilson coefficients in the SMEFT analysis, comparing the outcome of individual and global marginalised fits for the HL-LHC, LEP3, FCC-ee, LCF550, and LCF1000 projections. Note that since we absorb factors of Λ^{-2} in our definition of the Wilson coefficients c_i , these are dimensionful and have units of TeV^{-2} . These fits are carried out at $\mathcal{O}(\Lambda^{-2})$ in the EFT expansion; results accounting for quadratic corrections are presented in Sect. 7.2.2. One-loop RGE effects are included, and theoretical uncertainties are included in the ‘aggressive’ scenario; the impact of varying the latter is quantified in Sect. 7.2.3. In all cases, the bounds on the Wilson coefficients are provided at the reference scale of $\mu_0 = 10$ TeV.

The same information contained in Fig. 7.1 for the constraints of the different Wilson coefficients can be presented with alternative visualisations to facilitate its perusal. To this end, Figs. 7.2 and 7.3 display the same bounds as in Fig. 7.1, now in a spider-plot format, with the projected uncertainties in the bounds on the Wilson coefficients displayed as ratios to the HL-LHC expectations, for individual and global marginalised fits respectively. Specifically, we plot the metric $R_{\delta c_i}$, defined as the ratio between the magnitude of the 95% C.I. interval for a given EFT coefficient c_i obtained in a future collider fit, to



Figure 7.1: Results (length of the 95% C.I. bounds) of individual (triangles) and global marginalised (filled bars) fits to the projections for different future colliders for the $n_{\text{op}} = 61$ Wilson coefficients that compose our basis, displayed at the reference scale of $\mu_0 = 10$ TeV. These fits are carried out at the linear level, $\mathcal{O}(\Lambda^{-2})$, in the EFT expansion, account for RGE effects, and assume the aggressive scenario for theory uncertainties.

that of the same quantity in the HL-LHC fit:

$$R_{\delta c_i} = \frac{[c_i^{\min}, c_i^{\max}]^{95\% \text{ CI}} (\text{HL-LHC} + \text{future collider})}{[c_i^{\min}, c_i^{\max}]^{95\% \text{ CI}} (\text{HL-LHC})}, \quad i = 1, \dots, n_{\text{eff}}. \quad (7.16)$$

Therefore, in Figs. 7.2 and 7.3, for a given coefficient c_i , the smaller the value of $R_{\delta c_i}$ the more significant the impact of the new data. Note also the logarithmic scale of the radial axis, and that the Wilson coefficients are clustered in terms of the categories defined in Table 7.4. Furthermore, Fig. 7.4 shows the same results as in Fig. 7.1 presented as bounds on $\Lambda \equiv 1/\sqrt{c_i(\mu_0)}$, which can be interpreted as the mass scale being probed by the SMEFT analysis for Wilson coefficients of order unity at the reference scale $\mu_0 = 10 \text{ TeV}$.

To further streamline the interpretation of the results presented in Figs. 7.1–7.4, we also display in Fig. 7.5 the diagonal entries of the Fisher information matrix [317], evaluated at linear order in the EFT expansion for the LEP3, FCC-ee and LCF projections, in all cases with LEP and HL-LHC included. For the purposes of this Fisher information analysis, we cluster the input datasets into three groups: LEP, HL-LHC, and either LEP3, FCC-ee or LCF, in the latter case separating the different $\sqrt{s}$ runs. Each row of the table is normalised to 100. The larger the entry of these Fisher information matrices, the more dominant a specific dataset is in determining the corresponding operator in the case of individual fits. For instance, from Fig. 7.5 one observes that at the FCC-ee all four-quark operators are constrained predominantly by the Tera- Z observables, while for the LCF the impact of the Z -pole run for these operators is mostly subdominant and is instead driven by the energy runs above the Z -pole. As another example, at the LCF the two- and four-lepton operators are mostly determined by the $\sqrt{s} = 1 \text{ TeV}$ run, while at the FCC-ee they are mostly constrained by the runs at $\sqrt{s} = 240$ and 365 GeV .

We now comment on some of the main features observed from Figs. 7.1–7.5.

Individual versus marginalised fits. Comparing bounds obtained from individual (one-parameter) fits and marginalised fits in Figs. 7.1–7.4, a first observation is that the bounds obtained in the global marginalised case are in many cases much looser, often by orders of magnitude. This is the case for most of the four-fermion operators, both those including quarks and those with leptons, and also for some of the purely bosonic and two-fermion operators. These differences arise because, in the global fit (especially in the linear case), variations in one coefficient may be compensated with correlated variations of other coefficients while still leading to a satisfactory description of the experimental data, hence leading to a degradation of the bounds.

Individual and global marginalised bounds are instead similar for those Wilson coefficients that are weakly correlated in the global fit. Operators falling in this category include for example the self-coupling c_φ at the HL-LHC and LCF, which can directly be constrained from the di-Higgs production cross-section. Similar observations were made when comparing individual and marginalised fits in previous analysis [319].

Any realistic UV completion of the SM model will activate only a subset of dimension-six Wilson coefficients. At the same time, only very simple UV extensions would activate only a single SMEFT coefficient at the matching scale μ_0 . Hence the results presented in Figs. 7.1–7.4 need to be understood as bracketing what would be obtained for specific UV scenarios, where the bounds obtained will be more stringent than in the global fit, but looser as compared to the one-parameter analysis.

Analysis of individual bounds. The spider plot of Fig. 7.2 quantifies the expected improvement in the individual bounds for the Wilson coefficients that form our fitting basis in light of the future collider




Figure 7.2: Same as Fig. 7.1 now with the projected uncertainties in the individual bounds on the Wilson coefficients in a linear SMEFT fit with aggressive theory uncertainties displayed as ratios to the HL-LHC expectations, Eq. (7.16). We compare LEP3, FCC-ee, LCF550, and LCF1000. Note the logarithmic scale of the radial axis. The Wilson coefficients are clustered in terms of the categories defined in Table 7.4.

Ratio of Uncertainties to HL-LHC Baseline $\mathcal{O}(\Lambda^{-2})$, Marginalised




Figure 7.3: Same as Fig. 7.2 for the global marginalised bounds.




Figure 7.4: Same as Fig. 7.1 for the lower bounds on the mass scale $\Lambda \equiv c_i^{-1/2}(\mu_0)$, which can be interpreted as the characteristic NP mass scale being probed by the SMEFT analysis assuming Wilson coefficients of order unity. Results are evaluated at the reference energy scale of $\mu_0 = 10$ TeV.




Figure 7.5: The entries of the Fisher information matrix, evaluated at linear order in the EFT expansion, for the fits based on the LEP3, FCC-ee, LCF-550 and LCF-1000 projections (from left to right). For this analysis, we cluster the input datasets into three groups: LEP, HL-LHC, and either LEP3, FCC-ee, or LCF respectively, in the latter case separating the different $\sqrt{s}$ runs. Theoretical uncertainties are included in the aggressive scenario.

projections with respect to the HL-LHC baseline. Overall, the FCC-ee displays the best performance, with significant reduction of the bounds on the Wilson coefficients in all categories. For many directions in the SMEFT parameter space, the improvements of the FCC-ee as compared to the HL-LHC are better by a factor of at least 20. We note that these improvements also apply to operators, such as the two-light-two-heavy and four-heavy quark types, which are not directly accessed at lepton colliders but to which the fit is sensitive through loop effects via RGE evolution.

One also finds that a small number of operators will not be significantly improved in any of the future colliders considered compared to the HL-LHC, such as the top and muon Yukawa couplings $c_{t\varphi}$ and $c_{\mu\varphi}$, for which a higher-energy proton-proton or muon collider would be needed, or the purely bosonic operators $c_{\varphi W}$ and $c_{\varphi B}$. In the specific case of the top Yukawa coupling, we remark that our HL-LHC projections are not based on targeted analyses, and that variations of the correlation model for Higgs measurements at the HL-LHC can modify these expectations.

The LCF reach dominates over the FCC-ee one for some of the operators involving leptons, since these benefit from an increased sensitivity at the higher energy runs at $\sqrt{s} = 550$ GeV and 1 TeV. The LCF results are instead less competitive than the FCC-ee ones for the two-fermion-bosonic operators and most of the four-quark operators. These operators are constrained via RGE running to the Z -pole observables and are more weakly constrained at the Giga- Z and $\sqrt{s} = 250$ GeV runs of the LCF than at the Tera- Z run of the FCC-ee (and to a lesser extent of LEP3). We note in particular that the Tera- Z run at the FCC-ee is scheduled to have a luminosity almost 70 times larger than the Giga- Z and $\sqrt{s} = 250$ GeV runs combined at the LCF.

When considering individual bounds for the Wilson coefficients, LEP3 offers a competitive performance for most operators considered. The degradation of the bounds as compared to FCC-ee is typically around a factor two. As in the case of FCC-ee, individual fits at LEP3 constrain two-lepton-two-heavy quark operators even in the absence of the $e^+e^- \rightarrow t\bar{t}$ run through RGE evolution to the lower $\sqrt{s}$ observables.

Analysis of global marginalised bounds. When comparing the improvement in the bounds to the HL-LHC baseline for the global marginalised fits in Fig. 7.3, one observes qualitative differences as compared to the corresponding results at the individual fit level of Fig. 7.2. First, an inversion of the relative hierarchy in reach for the various colliders: while for individual fits FCC-ee is typically superior with a handful of exceptions, in the global marginalised fits the best performance is achieved by the LCF, especially if the run at $\sqrt{s} = 1$ TeV is taken into account. Second, the improvements with respect to the HL-LHC that could be obtained at the FCC-ee are now moderate, specifically for the two-light-two-heavy four quark operators. Third, LEP3, which was competitive with the other colliders in the individual fit, now leads to the worst constraints for all the operators, with, in many cases only a relatively small improvement compared to the HL-LHC baseline. Finally, for the LCF, the qualitative pattern of improvement as compared to HL-LHC remains similar in individual and in global marginalised fits, at odds with the FCC-ee case and LEP3.

The different qualitative behaviour displayed in Fig. 7.3 as compared to Fig. 7.2 may be understood as follows. For individual fits, FCC-ee excels thanks to its unprecedented precision of the Z -pole observables, to which most operators flow via RGE. In the global fit, however, these same Z -pole constraints are largely weakened by correlations, whereas the LCFs larger variety of observables helps break these correlations resulting in tighter marginalised constraints despite its lower overall sensitivity. Similarly, LEP3 can also probe with high precision the Z -pole observables, leading to strong indirect constraints

in the individual fit, but its lack of runs above $\sqrt{s} = 230$ GeV prevents competitive bounds in the marginalised fit.

The results of Figs. 7.2 and 7.3 may appear at first inspection to provide a contradictory message. However, this is not the case: any SMEFT interpretation depends, by construction, on the assumptions made on the operator basis. Confronting Fig. 7.3 with Fig. 7.2 therefore illustrates the two limiting cases: all operators are assumed to have the same prior relevance, or only one of them is assumed to fully dominate, in both cases at the matching scale μ_0 .

NP mass reach. When presented in terms of lower bounds in $\Lambda \equiv c_i^{-1/2}(\mu_0)$, as done in Fig. 7.4, the results of the individual and global marginalised fits can be interpreted in terms of the characteristic mass scale of UV physics being probed under the assumption that the dimensionless couplings satisfy $g_{UV} \sim 1$ at μ_0 (taken to be the same as the matching scale). From the results of individual fits, we find that future lepton colliders can probe values up to Λ around $\mathcal{O}(10)$ TeV, for instance FCC-ee reaches $\Lambda \sim 30$ TeV for two-fermion operators such as $c_{\varphi Q}^{(3)}$ and $c_{\varphi q}^{(3)}$.

Consistently with the trends displayed in the rest of the plots in this subsection, the highest mass reach for individual fits is achieved by the FCC-ee for the two-fermion and most of the four-quark operators, while the LCF typically dominates the reach for most of the four-fermion operators involving lepton fields, reaching for instance above $\Lambda \sim 100$ TeV for the four-lepton operators and for c_{Qe} .

When considering instead the bounds on the NP mass scale Λ obtained from the global fit after marginalisation, the reach on this UV mass scale is markedly reduced, with LCF displaying a superior performance. For instance, LCF reaches $\Lambda \sim 20$ TeV for the two-lepton-two-light-quark operators, 100 TeV for some of the four-lepton operators, and around 10 TeV for the two-lepton-two-heavy-quark operators. However, as discussed above, the bounds obtained in the global SMEFT fit are too conservative for any realistic UV model, and for specific UV completions, the reach in the mass scale Λ will lie between the bounds of the individual and global fits.

Correlation patterns. To conclude this discussion, we display in Fig. 7.6 the correlation matrix for the Wilson coefficients in a global SMEFT fit carried out at the linear level to LEP, (HL-)LHC, and FCC-ee projections. The numerical value of the correlation is shown only for those entries of the matrix satisfying $|\rho| > 0.2$. It is clear that the fit parameters are characterised by a complex correlation pattern, although some clear patterns exist. First of all, we note how operators within the same operator class are in general strongly correlated, as can be seen by inspecting for example the two-light-two-heavy operators, the four-heavy or the two-lepton-two-quark operators in Fig. 7.6. Beyond this, we note that cross-correlations between different operator classes also exist, an effect which is enhanced, in particular, through RG mixing effects. Obvious examples here include the four-heavy operators and the operators entering the EWPOs, for which we observe a strong correlation of $|\rho| > 0.7$. This is in stark contrast to the milder correlation pattern observed previously in the analysis of Ref. [330] where RGE effects were not considered.

7.2.2 Impact of quadratic EFT corrections

The results presented in Sect. 7.2.1 are based on EFT calculations truncated at the linear order in the EFT expansion. Quadratic corrections arising at $\mathcal{O}(\Lambda^{-4})$ from the squares of dimension-six amplitudes can be comparable, or in some cases even dominant, to the linear ones for certain observables and values of the Wilson coefficients. In order to quantify the stability of our results with respect to the inclusion of $\mathcal{O}(\Lambda^{-4})$ corrections, here we present similar results as those of Sect. 7.2.1 now keeping the quadratic


Correlation: FCC-ee, $\mathcal{O}(\Lambda^{-2})$

Figure 7.6: The correlation matrix in the space of Wilson coefficients in a global SMEFT fit carried out at the linear level to LEP, (HL-)LHC, and FCC-ee projections. The numerical value of the correlation is shown only for those entries of the matrix satisfying $|\rho| > 0.2$.

EFT corrections to the cross sections arising from these squares. All other ingredients of the global SMEFT fit are kept unchanged.

In order to assess the impact of quadratic corrections, we carried out two comparisons. First, Fig. 7.7 compares the global marginalised bounds on the Wilson coefficients obtained at the linear and quadratic level for each of the future colliders considered. Second, Fig. 7.8 compares the same marginalised bounds on the quadratic global EFT fit with the analogous results obtained in individual quadratic fits, where only one Wilson coefficient is assumed to be non-zero at the reference scale $\mu_0 = 10$ TeV. We now detail some observations that can be derived from Figs. 7.7 and 7.8.

Linear versus quadratic fits. We first note that for some operators, once quadratic corrections are included, the HL-LHC projections already saturate the sensitivity, with no or minor subsequent improvements from the lepton colliders. This is the case especially for the two-light-two-heavy quark operators.

Moreover, for some operators, one finds large differences in the bounds obtained in the linear and the quadratic EFT level, with the latter resulting in more stringent constraints by up to two orders of magnitude, for instance for most of the four-fermion operators. On the other hand, quadratic and linear bounds are similar for certain Wilson coefficients such as the Higgs self-coupling, the Yukawa operators, or the triple gauge coupling c_{WWW} . In the case of the LCF projections, differences between linear and quadratic bounds are smaller for operators involving a lepton bilinear, which benefit from the constraints from difermion production processes at high $\sqrt{s}$ values.

In general, from Fig. 7.7 one concludes that quadratic EFT corrections at the marginalised level can be sizeable for several of the Wilson coefficients included in the global fit, and therefore must be accounted for. They lead to phenomenologically relevant impact especially within the four-fermion operators and some of the two-fermion ones. Quadratic EFT corrections can, however, be safely neglected at future colliders for the Higgs self-coupling, the triple gauge coupling, and for the Yukawa couplings (except for the top Yukawa at the FCC-ee and LEP3).

Individual versus marginalised fits. We now consider the results shown in Fig. 7.8, which compares the bounds obtained from individual and marginalised analyses when quadratic EFT corrections are taken into account. The linear EFT fit results discussed in the previous section displayed large variations between the bounds on the Wilson coefficients obtained in the individual and global marginalised fits. The same trend is obtained in the case of quadratic fits.

As in the case of the linear fits, also for the quadratic ones the overall pattern of improvement across colliders as compared to the HL-LHC baseline varies dramatically between individual and global fits, as illustrated in particular by the two-fermion-two-bosonic operators. Concerning the individual quadratic bounds, the best performance is again obtained by the FCC-ee, with LCF being comparable or better for operators involving lepton bilinears. These plots illustrate some of the main findings that we just discussed: for individual quadratic fits, FCC-ee shows the best performance except for operators involving lepton bilinears and c_{WWW} . For global marginalised fits at the quadratic level, LCF leads to the largest reduction, and, finally, the benefits of the Tera- Z run are markedly reduced in the global fit.

Flat directions A major challenge in performing global SMEFT fits is the presence of flat or quasi-flat directions in the parameter space, where different operators lead to degenerate effects on the measured observables. To systematically resolve these flat directions and ensure a robust fit, we employ several complementary strategies. First, we redefine our fitting basis by using specific linear combinations of Wilson coefficients, rather than the raw operators themselves (see, for example, the coefficients



Figure 7.7: Same as Fig. 7.1, comparing the marginalised bounds obtained in global fits carried out at the linear level in the EFT expansion (filled bars) and those from the quadratic level fits (triangles). Theoretical uncertainties are included in the ‘aggressive’ scenario.




Figure 7.8: Same as Fig. 7.1 for quadratic fits carried out at $\mathcal{O}(\Lambda^{-4})$ in the EFT expansion. For each of the colliders considered, the marginalised bounds from the global fit (filled bars) are compared with the individual bounds (triangles). Theoretical uncertainties are included in the ‘aggressive’ scenario.

indicated with an asterisk in Table 7.2, or the explicit linear combinations defined in Table 7.3). Second, the specific flavour symmetry assumptions adopted in our analysis naturally restrict the operator count and help lift parameter degeneracies. Third, the inclusion of RGE evolution introduces operator mixing across different energy scales, which provides indirect sensitivity to otherwise unconstrained coefficients and helps break degeneracies. Finally, the inclusion of quadratic $\mathcal{O}(1/\Lambda^4)$ EFT effects in our theoretical predictions further breaks residual correlations that would otherwise remain perfectly flat if the fit were truncated strictly at the linear level.

7.2.3 Impact of theoretical uncertainties

As demonstrated in Refs. [328,337], theoretical uncertainties can significantly degrade the physics reach of future lepton colliders, especially for the high-precision Z -pole measurements, but also for Higgs and difermion production. Here we quantify the role played by theory uncertainties in our SMEFT interpretations of future lepton collider data. Fig. 7.9 presents a similar comparison for the lower bound on the NP mass scale $\Lambda = c_i^{-1/2}(\mu_0)$ as in Fig. 7.4 for the global marginalised bounds in the linear EFT fit, now for the results in the four scenarios for theoretical uncertainties discussed in Sect. 7.1.5: current, conservative, aggressive, and ideal. Notice that in these comparisons the x -axis has a linear scale, and its range has been adjusted to match the typical variations in the different groups of Wilson coefficients.

From Fig. 7.9 one finds that theoretical uncertainties reduce the reach in the mass scale Λ by an amount that depends strongly on the specific operators being probed, the collider, and the scenario for the expected reduction of theoretical uncertainties assumed. For instance, for the purely bosonic operator $c_{\varphi WB}$, the projected bounds on Λ improve from 7.5 TeV to almost 12 TeV (10 TeV to 15 TeV) at the FCC-ee (LCF550), when moving from the current to the ideal theory scenario. For other operators, theoretical uncertainties barely have an impact, such as for the Higgs self-coupling c_φ , or the triple gauge coupling c_{WWW} . In most cases, differences between the ‘conservative’ and ‘ideal’ scenario are found to be at the 10% to 20% level for the global marginalised fit results.

In Fig. 7.10 we present a similar comparison as in Fig. 7.9 for the case of individual fits. There we find, instead, that the impact of theoretical uncertainties is rather more significant than for the global marginalised analysis. The observation that stands out the most is how different scenarios for the theory uncertainties impact the individual bounds at FCC-ee and LEP3 for operators that are sensitive to the Z -pole observables. For example, the $c_{\varphi q}^3$ coefficient shows variations up to a factor 13 at the FCC-ee depending on whether the current or ideal theory scenario is considered. By contrast, the effect of the different scenarios is less pronounced at LCF, because FCC-ee is significantly more sensitive to Z -pole observables and is therefore more strongly affected by the corresponding theoretical uncertainties.




Figure 7.9: Same as Fig. 7.4 for the marginalised lower bounds on the mass scale $\Lambda = c_i^{-1/2}(\mu_0)$ in the global linear fits at the future colliders considered, comparing the results in the four scenarios for theoretical uncertainties at future lepton colliders: current, conservative, aggressive, and ideal. Note that now the x -axis scale is linear and varies for each group of Wilson coefficients.




Figure 7.10: Same as Fig. 7.9 now for individual bounds on the Wilson coefficients.

Conclusion

“There are other Annapurnas in the lives of men.”

Maurice Herzog, *french alpinist, first ascensionist of an 8000m peak (Annapurna).*

The research presented in this thesis was driven by a paradigm shift in the high-energy physics program: the transition from an era of discovery to an era of precision. While the Standard Model remains an exceptionally description of fundamental interactions up to the TeV scale, its inability to account for phenomena such as dark matter and the baryon asymmetry hints that it is a low-energy approximation of a more general ultraviolet theory. In the absence of direct signals of beyond the Standard Model physics at the LHC, the search for new fundamental laws has become intrinsically linked to our ability to predict Standard Model processes with sub-percent accuracy. The central focus of this thesis is the calculation of higher order corrections to physical observables.

To achieve this level of theoretical precision, a major bottleneck in state-of-the-art calculations is the evaluation of complex, multi-scale Feynman integrals. Prior to this work, public tools such as DIFFEXP, were available for evaluating these integrals via series expansion, however being restricted to real-valued internal masses. In this thesis, we developed and published the open-source MATHEMATICA package SEASYDE. SEASYDE extends the series expansion approach to the complex domain, allowing for a robust treatment of intermediate unstable particles like W and Z bosons within the complex mass scheme. Furthermore, we integrated this solver into the OCEANN framework, establishing a highly automated pipeline capable of handling the generation, reduction, and numerical evaluation of two-loop amplitudes.

Using this machinery, we achieved a significant milestone in precision phenomenology: the computation of the exact two-loop mixed QCD-EW amplitudes for both Neutral-Current and Charged-Current Drell-Yan processes. Previously, phenomenological studies of these mixed corrections relied heavily on the pole approximation or factorized approaches. Our exact computation overcomes these limitations. By implementing the complex mass scheme and retaining the exact lepton mass dependence to regulate final-state collinear singularities, we computed the two-loop virtual amplitudes across the entire phase space without relying on resonant expansions. As a first step towards full two-loop EW calculations, we also successfully isolated and subtracted the infrared pole structure for the two-loop QED corrections to Neutral-Current Drell-Yan.

By embedding the two-loop QCD-EW amplitudes for Neutral-Current Drell-Yan within the MATRIX framework, we produced state-of-the-art, fully differential fiducial cross-sections in both the resonant and high invariant mass region of the phase space. We showed that the exact mixed QCD-EW corrections become extremely relevant in the high invariant mass region where their impact could reach a few percents.

Finally, we bridged the gap between Standard Model predictions and the active search for new physics through the framework of the SMEFT. Moving beyond isolated measurements, we performed a comprehensive global SMEFT analysis to project the indirect discovery reach of future electron-positron colliders (such as FCC-ee, LEP3, and LCF). Our global fit analysis included 2-lepton-2-quark operators, quadratic effects, and accounted for Renormalization Group Equation effects.

Looking ahead, the demand for theoretical precision will only intensify. The upcoming High-Luminosity LHC program will deliver an unprecedented volume of data, reducing statistical uncertainties for a vast array of standard-candle observables down to the per-mille level. To ensure that theoretical uncertainties do not become the primary limiting factor in the indirect search for BSM physics, several critical mathematical and computational challenges must be addressed by the community in the coming years: full two-loop EW corrections for $2 \rightarrow 2$ processes and mixed two-loop corrections for $2 \rightarrow 3$ ones. As experimental precision steadily increases, the computation of the full NNLO EW corrections for $2 \rightarrow 2$ scattering processes will become an unavoidable phenomenological requirement. As highlighted by the benchmark studies performed in this thesis, transitioning to the full electroweak sector induces an explosive growth in complexity. This entails managing thousands of Feynman diagrams and evaluating non-planar topologies with multiple internal mass scales, which leads to systems of differential equations reaching hundreds of megabytes in size. Overcoming this bottleneck will require advancing rational reconstruction techniques and casting the system in canonical bases to simplify the integration of these massive systems. Beyond $2 \rightarrow 2$ scattering, extending exact NNLO QCD-EW calculations to $2 \rightarrow 3$ processes, such as $q\bar{q} \rightarrow \mu^+\mu^-g$, represents the next major theoretical frontier. Processes like vector boson plus jet production are crucial backgrounds for numerous BSM searches and play a central role in precise detector calibration. Computing the two-loop virtual amplitudes for these reactions requires handling non-trivial loop integrals with up to five external legs and numerous mass scales. Evaluating these five-point functions currently pushes the absolute limits of modern reduction algorithms, demanding further fundamental developments in the analytical and numerical evaluation of multi-scale Feynman integrals.

In conclusion, the work presented in this thesis directly contributes to the ongoing effort to interrogate the Standard Model at unprecedented levels of accuracy. By mastering the computation of multi-loop amplitudes, developing automated tools for their evaluation, and translating them into rigorous phenomenological constraints, we continue to lay the necessary groundwork for the next great discovery in particle physics. The precision frontier is not merely a refinement of known physics, but it is the most powerful telescope we possess for discovering the unknown.

Acknowledgements

“Aaaaaaaaaahhhhhhhhhhhh”

Adam Ondra on Silence 9c, *Czech climber.*

The alarm rings. Its 2:00 AM. Next to me, I listen to the others getting out of bed. I slowly crawl out as well and start helping make the tea for breakfast. I am tired, and I am scared. When Cami asked me whether I wanted to join them to climb one of the most remote 4000 m peaks in the Alps, I said yes without thinking twice. Now, I start wondering if this was a good idea. Yesterday, we walked the entire day: 17 km and 1400 m of elevation gain with heavy backpacks across valleys and glaciers. When we finally arrived at the bivouac and looked at our goal, the Lauteraarhorn (4042 m), I felt intimidated. The mountain is huge, the nearest city is 7-8 hours walk away, and the route seems more technical than I expected. While I am deep in my fears, breakfast is ready. We eat bread with peanut butter and jam, drink our tea, finish preparing the last few things, and at 3:00 AM, we leave the hut. The moon is high in the sky and we almost don't need our headlamps. Despite being at 3000 m at 3:00 AM, the temperature is abnormally high: a few years ago, people used to do this approach with skis in August, now we are walking on rocks in t-shirts, another sign of global warming. The approach is long, but not difficult. We walk on rocks for a few hours until we reach a glacier. We put on crampons and start crossing it. However, the temperature is too high and the snow is not hard enough, so we decide to rope up. After crossing the glacier, we finally get to the bottom of the South couloir, the real start of our climb. The sun is rising. Its beautiful. It is just us, the mountains around us, and the silence. There is just enough time for a snack before we start climbing. The conditions of the couloir are far from perfect. There has been almost no night frost, so despite the early hour, the snow is soft. With a better freeze, one would be able to climb the 55 degree couloir in almost no time with crampons and ice axes. In these conditions, it becomes extremely tiring. At every step, you sink to your hips and have to fight for each meter. Following the leader is of course easier, so we decide to alternate. Despite this, I am already exhausted. We are almost at 4000 m, there is not much oxygen in the air, and if it was not for my friends, I would have already turned back. At 10:00 AM, we finally get to the top of the couloir at 3900 m. Time for a quick snack before the last 100 m on the south ridge. Now we are moving on rocks, a type of terrain where I feel more comfortable, so I hope we can speed up a little bit. We are six, so we rope up in teams of two and start simul-climbing. After a few meters, there is a tricky downclimb. Ale, Cami, Edo, and Leo decide to turn back because of the warm temperatures and the late hour. Ettore and I decide to continue. The climb is just a III grade, we move fast on the ridge, and at 11:30 AM we finally reach the top! Being on top of your first 4000 m peak is something you never forget. Its just you for kilometers. You feel the cold wind hitting your face and take in a panorama that not many other people have seen. In that moment, I felt so little in front of the Earth. We don't have much time. The descent is long, and the temperature is rising, meaning the snow will become softer and softer. We just have time to eat a snack, take a picture, sign the summit logbook, and then we start descending, trying to catch up with the others. The first part of the ridge goes smoothly, and in no time we are once again at the top of the couloir. Now the dangerous part begins. Due to the high temperature, the snow is unstable, and at every step, there

is a risk of an avalanche. Below us, we see the others and they are moving slower than expected. The descent is sketchy, and at every step, you have to make sure your feet are stable, otherwise you might start sliding down. Halfway through the couloir, we catch up with the others. A few moments later, an avalanche catches Leo. He slides for 30 meters before stopping on some rocks. Everybody holds their breath. He gets up and waves at us. I have never been so happy to see somebody being okay. We have to hurry up: its getting dangerous. Finally, at 2:00 PM, we arrive at the bottom of the couloir. The only obstacle between us and the end of the difficulties is crossing the glacier. The snow is soft now, so the risk of falling into a crevasse is real. We rope up. Leo and I are the more inexperienced ones, so we have to go in front. He has just been caught in an avalanche, so he asks me to go first. I accept. I start walking. At every step, I sink down to my knees. Until, at some point, my foot sinks deeper, and before I can realize what is happening, I feel the void below me. Leo does a wonderful job catching me, and I find myself hanging in a crevasse up to my shoulders. They pull me out, and we finally get to the rocks, the end of the difficulties. From here, it is just a long walk to the car, which we will reach only at 1:00 AM the next day.

I think the story of my first 4000 m peak is a perfect metaphor for what my PhD has been. A loooooong journey, filled with difficulties, obstacles, and unforeseen events, which pushed me to my limits on many occasions. At the same time, it brought me to beautiful places where very few people have been, and gave me the possibility to lay my eyes on breathtaking landscapes. And in a PhD, just as in the mountains, you always need people. To guide you, to pull you out of crevasses, to cheer you up in difficult moments, or simply to drink a beer with at the end of the day. Thats why I want to thank all the people who shared a part of this journey with me.

First and foremost, thanks to my supervisors. Thanks, Alessandro, for guiding me through my master's thesis first and for pushing me to pursue the PhD later. Thanks for the long discussions, for your patience in explaining things to me multiple times, and for giving me the freedom to explore ideas. Thanks, Fabio, for giving me the possibility to do my PhD in Louvain, which has been the place I felt at home for a few years. Thanks for the insightful discussions, for always giving the right idea at the right time, and for putting me in contact with many beautiful people during these years. Thanks to Gauthier, Jonas, Philippe and Stefano, my PhD committee. Thanks for your comments on the thesis, I really appreciated you taking the time to go so in-depth on it. Thanks to Simone for being my older brother on this journey. Thanks for teaching me everything, from running slurm on the cluster to how qt-subtraction works (multiple times, actually), for answering all the stupid questions I was scared of asking the boss, and, most importantly, for the countless board game nights. Thanks, Michele, for the support, the discussions, the racchettoni matches, and the Seaside games. Having someone to share the journey with, including the highs and the lows, was fundamental to me. Thanks to Narayan and Roberto for teaching me many technicalities about Feynman integrals, renormalization, and for always being supportive.

Thanks to my family: Barbara, Riccardo, Sofia, and Ynse. Thanks for the continuous support, for trying (multiple times, without a lot of success) to understand what a loop integral is, and for pushing me to pursue what I liked most without forcing me, even though you did not fully understand my choices.

Thanks to my Louvain family: Daniele, Emile, Federico, Francisco, Jonathan, Leonardo, Luca, Marco, Marco, Paola, Ricardo, Simone, Stavros, Valentin, and Ze-Qiang. I loved every moment I spent with you, the lunches and the nights, the mitragliettes and the calcetto in the Beer Bar. You were the reason I could call Louvain home for all these years. Thanks to the friends I made in Milan. Thanks, Andrea and Davide, for the laughs in the office, Eva for the ski touring, and Sabina for the shottini tequila, sale e limone like we were 18!


Thanks to my lifetime friends. Even though we were apart, I always felt at home when I came back to Milan. Thanks, Mattia and Martina, giving the best man speech at the kebab shop will always be one of the proudest moments of my life! Thanks, Federico, for visiting me in Belgium I do not know how many times. You were a reference point for me and helped me so much in difficult times. Thanks, Francesco, for being at my side since middle school, for the Paleo nights, and for shooting naked in the woods. Thanks, Greta and Marta, for all the adventures, for crossing Europe in a Flixbus for me, for the Scottish holidays, and, more generally, for all the laughs in front of a wine glass. Thanks, Andrea, Andrea, Aurora, Camilla, Eleonora, Elisa, and Michela, for always being the fixed point when everything else was changing.

Thanks, Nicola, for helping me express my full potential. I know you wanted to be in the acknowledgments for the Nobel Prize, you will have to settle for the PhD ones!

Thanks to my climbing friends. Thanks to Mark for spending months and months in Louvain and for making me discover climbing. Thanks to Camila and *mon pain suisse*, Matteo, for the crazy stupid climbing/biking/hiking holidays in Verdon and in the Dolomites. I actually cannot wait to see what our next adventure will be! Thanks, Elisa, Luca, and Michelle, for the countless climbing nights at Solid. Thanks to my CAI friends. In particular, to Leonardo for pulling me out of that crevasse, but also for all the other adventures, including the Dent Blanche attempt, which, to date, is one of the best experiences I have had in the mountains despite the fact we did not manage to reach the top. Thanks to Pietro for howling on top of the Grand Capucin. Thanks to Alessandro, Anna, Arsh, Camilla, Cristina, Edoardo, Ettore, Francesco, Gilberto, Jennifer, Lorenzo, Marta, Matteo, Maria, Michele, Mimma, Paolo, and Xenia for the climbs, laughs, and beers we had and will have in the future! Thanks, Alessio and Eleonora, for the beautiful adventure on the Alta Via. I will never forget those beautiful sceneries.

Last but not least, thanks to Luciana, the person who, in this last year, pushed me through hard times even when I wanted to give up. Thanks for all our little climbing adventures in the mountains and at the sea. For the laughs, the picnic on the lake, for the good movies you always choose, I feel safe just being myself when I am with you. I could not be any luckier to have met you. Thanks *bro*!

Thanks again to all of you, whether you shared the whole journey or just a part of it. In research, just as in the mountains, a reliable team is fundamental to reaching the summit, and I could not be happier with the team I found. This specific expedition is now complete, and I cannot wait to see what the next peak will be!



Photos by Camilla Ielmini.

Renormalisation relations

We report here the basic definitions that we have followed for the semi-analytical computation of the counterterms, needed to obtain the complete two-loop renormalised amplitude.

A.1 Electric charge renormalisation within the $\overline{MS}$ scheme

The renormalisation of the electric charge has been carried out following the same scheme for both the diagrammatic contributions and the subtraction operators. The relation between the bare electric charge and the renormalised one is defined up to the second order in α by introducing the counterterms

$$e_0 = e + \delta e = e Z_e = e \left(1 + \delta Z_e^{(1)} + \delta Z_e^{(2)} + \mathcal{O}(\alpha^3) \right). \quad (\text{A.1})$$

Following the on-shell renormalisation scheme, the Ward Identity associated to the residual $U(1)_{em}$ gauge invariance ensures that the renormalisation of the electric charge follows entirely from the transverse photon-photon self-energy $\Sigma_T^A(q^2)$ [48, 79, 94], via the renormalisation condition on the $\gamma f f$ vertex that relates Z_e with the photon wavefunction counterterm $Z_{AA} = 1 + \delta Z_{AA}^{OS} = 1 + \delta Z_{AA}^{OS(1)} + \delta Z_{AA}^{OS(2)} + \dots$:

$$Z_e^{OS} = \frac{1}{\sqrt{Z_{AA}^{OS}}} = \frac{1}{\sqrt{1 + \delta Z_{AA}^{OS(1)} + \delta Z_{AA}^{OS(2)} + \dots}}, \quad (\text{A.2})$$

with

$$\delta Z_{AA}^{OS} = - \frac{\partial \Sigma_T^A(q^2)}{\partial q^2} \Big|_{q^2=0}. \quad (\text{A.3})$$

In the $\overline{MS}$ scheme, only the divergent terms of the self-energies are kept in the definitions of the counterterms, and an additional dependence on the renormalisation scale μ_R and on the regularisation 't Hooft scale μ is introduced together with an overall factor $S_\epsilon = (4\pi e^{-\gamma_E})^\epsilon$:

$$\begin{aligned} \delta Z_e^{\overline{MS}(1)} &= S_\epsilon \left(\frac{\mu^2}{\mu_R^2} \right)^\epsilon \left[-\frac{1}{2} \delta Z_{AA}^{OS(1)} \Big|_{\text{poles}} \right], \\ \delta Z_e^{\overline{MS}(2)} &= S_\epsilon^2 \left(\frac{\mu^2}{\mu_R^2} \right)^{2\epsilon} \left[-\frac{1}{2} \delta Z_{AA}^{OS(2)} \Big|_{\text{poles}} + \frac{3}{8} \left(\delta Z_{AA}^{OS(1)} \Big|_{\text{poles}} \right)^2 \right]. \end{aligned} \quad (\text{A.4})$$

We have checked our results to be in agreement with an arbitrary number of digits with the expressions provided in [276], that we report here for completeness after switching to our notation:

$$\begin{aligned}\delta Z_e^{\overline{MS}(1)} &= \frac{\alpha(\mu_R)}{4\pi} S_\epsilon \left(\frac{\mu^2}{\mu_R^2} \right)^\epsilon \left[\frac{2}{3\epsilon} \sum_k N_k Q_k^2 \right], \\ \delta Z_e^{\overline{MS}(2)} &= \left(\frac{\alpha(\mu_R)}{4\pi} \right)^2 S_\epsilon^2 \left(\frac{\mu^2}{\mu_R^2} \right)^{2\epsilon} \left[\frac{2}{3\epsilon^2} \left(\sum_k N_k Q_k^2 \right)^2 + \frac{1}{\epsilon} \sum_k N_k Q_k^4 \right].\end{aligned}\tag{A.5}$$

In our computation, both the UV renormalisation procedure and the IR subtraction of divergences have been realized fixing for convenience $\mu = \mu_R$, resulting in an expression for δZ_e completely independent of any arbitrary scale.

A.2 Fermion mass and wave-function on-shell renormalisation

The fermion fields and masses are renormalised up to $\mathcal{O}(\alpha^2)$ with the wave-function renormalisation constants:

$$\begin{aligned}m_{f0} &= m_f + \delta m_f = m_f + \delta m_f^{(1)} + \mathcal{O}(\alpha^2), \\ f_0 &= (Z_f)^{\frac{1}{2}} f = \left[1 + \frac{1}{2} \delta Z_f^{(1)} + \frac{1}{2} \delta Z_f^{(2)} - \frac{1}{8} \left(\delta Z_f^{(1)} \right)^2 + \mathcal{O}(\alpha^3) \right] f.\end{aligned}\tag{A.6}$$

The counterterm expressions has been computed with a semi-analytical evaluation fixing the on-shell renormalisation conditions satisfied from the fermion two-point vertex function. Following the same notation of [27, 48] with $\Sigma_L^{f\bar{f}} = \Sigma_R^{f\bar{f}} = \Sigma^{f\bar{f}}$, we have:

$$\begin{aligned}\delta m_f^{(1)} &= m_f \left[\Sigma^{f\bar{f}(1)}(m_f^2) + \Sigma_S^{f\bar{f}(1)}(m_f^2) \right], \\ \delta Z_f^{(1)} &= -\Sigma^{f\bar{f}(1)}(m_f^2) - 2m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma^{f\bar{f}(1)}(p^2) + \Sigma_S^{f\bar{f}(1)}(p^2) \right] \Big|_{p^2=m_f^2}, \\ \delta Z_f^{(2)} &= -\Sigma^{f\bar{f}(2)}(m_f^2) - 2m_f^2 \frac{\partial}{\partial p^2} \left[\Sigma^{f\bar{f}(2)}(p^2) + \Sigma_S^{f\bar{f}(2)}(p^2) \right] \Big|_{p^2=m_f^2}.\end{aligned}\tag{A.7}$$

Our results for the fermion counterterms are in agreement with the analytic formula present in [276] after including the self-energy dependence on the different internal masses.

Photon vacuum polarization

The change of renormalisation scheme for the QED coupling can be expressed with a finite renormalisation between the on-shell one and $\overline{MS}$ schemes. Following the approach described in [79], we define the photon vacuum polarization function from the self-energy expression:

$$\Sigma_T^{AA}(q^2) = - (4\pi\alpha) q^2 \Pi_{\gamma\gamma}(q^2) . \quad (\text{B.1})$$

Indicating with $\bar{\Pi}^{(i)}(0)$ only the ϵ -divergent contribution of $\Pi^{(i)}(0)$ we therefore link the photon vacuum polarization function with our definitions of wave-function counterterms:

$$\begin{aligned} \delta Z_{AA}^{OS} = 4\pi\alpha \Pi_{\gamma\gamma}(0) &\implies \delta Z_{AA}^{OS(i)} = 4\pi\alpha^i \Pi_{\gamma\gamma}^{(i)}(0) , \\ \delta Z_{AA}^{\overline{MS}} = 4\pi\alpha_{\overline{MS}} \bar{\Pi}_{\gamma\gamma}(0) &\implies \delta Z_{AA}^{\overline{MS}(i)} = 4\pi(\alpha_{\overline{MS}})^i S_\epsilon^i \bar{\Pi}_{\gamma\gamma}^{(i)}(0) . \end{aligned} \quad (\text{B.2})$$

The two expressions are related by the definition of the bare coupling

$$\alpha_0 = \alpha Z_e^2 = \alpha (Z_{AA}^{OS})^{-1} = \alpha_{\overline{MS}} (Z_{AA}^{\overline{MS}})^{-1} , \quad (\text{B.3})$$

that we can expand up to the second order in α and/or $\alpha_{\overline{MS}}$ (since $\mathcal{O}(\alpha) = \mathcal{O}(\alpha_{\overline{MS}})$):

$$\begin{aligned} \alpha_{\overline{MS}} &= \alpha - 4\pi\alpha\alpha_{\overline{MS}} \left(\Pi_{\gamma\gamma}^{(1)}(0) - \bar{\Pi}_{\gamma\gamma}^{(1)}(0) \right) \\ &\quad - 4\pi\alpha(\alpha_{\overline{MS}})^2 \left(\frac{\alpha}{\alpha_{\overline{MS}}} \Pi_{\gamma\gamma}^{(2)}(0) - \bar{\Pi}_{\gamma\gamma}^{(2)}(0) \right) + \mathcal{O}(\alpha^4) \\ &= \alpha \left(1 - 4\pi\alpha_{\overline{MS}} \hat{\Pi}_{\gamma\gamma}^{(1)}(0) - 4\pi(\alpha_{\overline{MS}})^2 \hat{\Pi}_{\gamma\gamma}^{(2)}(0) + \mathcal{O}(\alpha^3) \right) , \end{aligned} \quad (\text{B.4})$$

where the factors S_ϵ^i have been removed by rescaling the t'Hooft parameter in $\Pi_{\gamma\gamma}(0)$ with $\mu \rightarrow \sqrt{e^{\gamma_E}/4\pi} \mu$. In the last row we exploited the fact that $\alpha/\alpha_{\overline{MS}} = 1 + \mathcal{O}(\alpha)$ and we defined $\hat{\Pi}_{\gamma\gamma}^{(i)}(0)$ as the finite remainder from the difference $\Pi_{\gamma\gamma}^{(i)}(0) - \bar{\Pi}_{\gamma\gamma}^{(i)}(0)$. Equation B.4 directly leads to the definition of $\Delta\alpha_{\overline{MS}}(m_Z^2)$ after setting $\mu_R = m_Z$. The photon vacuum polarisation can be decomposed as:

$$\begin{aligned} \Pi_{\gamma\gamma}(0) &= \Pi_{\gamma\gamma}^{(\text{lep})}(0) + \Pi_{\gamma\gamma}^{(\text{top})}(0) + \Pi_{\gamma\gamma}^{(5)}(0) \\ &= \Pi_{\gamma\gamma}^{(\text{lep})}(0) + \Pi_{\gamma\gamma}^{(\text{top})}(0) + \text{Re}(\Pi_{\gamma\gamma}^{(5)}(m_Z^2)) - \frac{\Delta\alpha_{\text{had}}(m_Z^2)}{4\pi\alpha} , \end{aligned} \quad (\text{B.5})$$

where (lep) and (top) are the perturbative contributions from the leptons and the massive top quarks, (5) is the contribution from the five light quarks and $\Delta\alpha_{\text{had}}(m_Z^2) = 4\pi\alpha \left(\text{Re}(\Pi_{\gamma\gamma}^{(5)}(m_Z^2)) - \Pi_{\gamma\gamma}^{(5)}(0) \right)$ is derived via a dispersion relation from the experimental cross section of $e^+e^- \rightarrow \text{hadrons}$. In conclusion

we obtain:

$$\begin{aligned}
 \Delta\alpha_{\overline{MS}}(m_Z^2) &= \Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2) + \Delta\alpha_{\overline{MS}}^{(2)}(m_Z^2) + \mathcal{O}(\alpha^3) = \\
 &= -4\pi\alpha_{\overline{MS}} \left[\hat{\Pi}_{\gamma\gamma}^{(\text{lep})^{(1)}}(0) + \hat{\Pi}_{\gamma\gamma}^{(\text{top})^{(1)}}(0) + \text{Re}(\hat{\Pi}_{\gamma\gamma}^{(5)^{(1)}}(m_Z^2)) - \frac{\Delta\alpha_{\text{had}}(m_Z^2)}{4\pi\alpha} \right] \\
 &\quad - 4\pi(\alpha_{\overline{MS}})^2 \left[\hat{\Pi}_{\gamma\gamma}^{(\text{lep})^{(2)}}(0) + \hat{\Pi}_{\gamma\gamma}^{(\text{top})^{(2)}}(0) + \text{Re}(\hat{\Pi}_{\gamma\gamma}^{(5)^{(2)}}(m_Z^2)) \right] + \mathcal{O}(\alpha^3) .
 \end{aligned} \tag{B.6}$$

The term $\Delta\alpha_{\text{had}}(m_Z^2)$ starts at $\mathcal{O}(\alpha)$ and takes part to the definition of $\Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2)$, thus affecting the change of scheme for the virtual amplitude starting from NLO. The term $\Delta\alpha_{\text{had}}(m_Z^2)$ starts at $\mathcal{O}(\alpha)$ and takes part to the definition of $\Delta\alpha_{\overline{MS}}^{(1)}(m_Z^2)$, thus affecting the change of scheme for the virtual amplitude starting from NLO.

Bibliography

- [1] Tommaso Armadillo, Roberto Bonciani, Simone Devoto, Narayan Rana, and Alessandro Vicini, “Two-loop mixed QCD-EW corrections to neutral current Drell-Yan”, *JHEP*, vol. 05, pp. 072, 2022.
- [2] Tommaso Armadillo, Roberto Bonciani, Simone Devoto, Narayan Rana, and Alessandro Vicini, “Evaluation of Feynman integrals with arbitrary complex masses via series expansions”, *Comput. Phys. Commun.*, vol. 282, pp. 108545, 2023.
- [3] Tommaso Armadillo, Roberto Bonciani, Simone Devoto, Narayan Rana, and Alessandro Vicini, “Two-loop mixed QCD-EW corrections to charged current Drell-Yan”, *JHEP*, vol. 07, pp. 265, 2024.
- [4] Tommaso Armadillo, Roberto Bonciani, Luca Buonocore, Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, Narayan Rana, and Alessandro Vicini, “Mixed QCD-EW corrections to the neutral-current Drell-Yan process”, *JHEP*, vol. 07, pp. 141, 2025.
- [5] Tommaso Armadillo, “Evaluating Feynman integrals through differential equations and series expansions”, *Eur. Phys. J. ST*, vol. 234, no. 26, pp. 7989–8004, 2026.
- [6] Tommaso Armadillo, Simone Devoto, Michele Dradi, and Alessandro Vicini, “Towards the two-loop electroweak corrections to the Drell-Yan process: the infrared structure”, 2025.
- [7] Tommaso Armadillo, Eugenia Celada, Jaco ter Hoeve, Fabio Maltoni, Luca Mantani, Juan Rojo, Alejo N. Rossia, Simone Tentori, Marion O. A. Thomas, and Eleni Vryonidou, “New Physics Reach through Precision at Future Colliders: a Multi-Pronged Approach”, 2026.
- [8] Tommaso Armadillo, Simone Devoto, Michele Dradi, and Alessandro Vicini, “Towards the two-loop electroweak corrections to the Drell-Yan process: the complete fermionic contributions”, 2026.
- [9] S. Schael et al., “Precision electroweak measurements on the Z resonance”, *Phys. Rept.*, vol. 427, pp. 257–454, 2006.
- [10] Georges 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*, vol. 716, pp. 1–29, 2012.
- [11] Serguei Chatrchyan et al., “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”, *Phys. Lett. B*, vol. 716, pp. 30–61, 2012.
- [12] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons”, *Phys. Rev. Lett.*, vol. 13, pp. 321–323, 1964.
- [13] Peter W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons”, *Phys. Rev. Lett.*, vol. 13, pp. 508–509, 1964.

- [14] Peter W. Higgs, “Broken symmetries, massless particles and gauge fields”, *Phys. Lett.*, vol. 12, pp. 132–133, 1964.
- [15] Peter W. Higgs, “Spontaneous Symmetry Breakdown without Massless Bosons”, *Phys. Rev.*, vol. 145, pp. 1156–1163, 1966.
- [16] P. Azzi et al., “Report from Working Group 1: Standard Model Physics at the HL-LHC and HE-LHC”, *CERN Yellow Rep. Monogr.*, vol. 7, pp. 1–220, 2019.
- [17] M. Cepeda et al., “Report from Working Group 2: Higgs Physics at the HL-LHC and HE-LHC”, *CERN Yellow Rep. Monogr.*, vol. 7, pp. 221–584, 2019.
- [18] Ansgar Denner and Stefano Pozzorini, “One loop leading logarithms in electroweak radiative corrections. 1. Results”, *Eur. Phys. J. C*, vol. 18, pp. 461–480, 2001.
- [19] D. de Florian et al., “Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector”, *CERN Yellow Rep. Monogr.*, vol. 2, pp. 1–869, 2017.
- [20] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, *JHEP*, vol. 07, pp. 079, 2014.
- [21] Gudrun Heinrich, “Collider Physics at the Precision Frontier”, *Phys. Rept.*, vol. 922, pp. 1–69, 2021.
- [22] Stefan Weinzierl, “The Art of computing loop integrals”, *Fields Inst. Commun.*, vol. 50, pp. 345–395, 2007.
- [23] Johannes M. Henn, “Lectures on differential equations for Feynman integrals”, *J. Phys. A*, vol. 48, pp. 153001, 2015.
- [24] Michael E. Peskin and Daniel V. Schroeder, *An Introduction to quantum field theory*, Addison-Wesley, Reading, USA, 1995.
- [25] Steven Weinberg, *The Quantum theory of fields. Vol. 1: Foundations*, Cambridge University Press, 6 2005.
- [26] Matthew D. Schwartz, *Quantum Field Theory and the Standard Model*, Cambridge University Press, 3 2014.
- [27] Ansgar Denner and Stefan Dittmaier, “Electroweak Radiative Corrections for Collider Physics”, *Phys. Rept.*, vol. 864, pp. 1–163, 2020.
- [28] H. Fritzsch, Murray Gell-Mann, and H. Leutwyler, “Advantages of the Color Octet Gluon Picture”, *Phys. Lett. B*, vol. 47, pp. 365–368, 1973.
- [29] D. J. Gross and Frank Wilczek, “Asymptotically Free Gauge Theories - I”, *Phys. Rev. D*, vol. 8, pp. 3633–3652, 1973.
- [30] D. J. Gross and Frank Wilczek, “Asymptotically Free Gauge Theories 2”, *Phys. Rev. D*, vol. 9, pp. 980–993, 1974.

- [31] H. David Politzer, “Reliable Perturbative Results for Strong Interactions?”, *Phys. Rev. Lett.*, vol. 30, pp. 1346–1349, 1973.
- [32] S. L. Glashow, “Partial Symmetries of Weak Interactions”, *Nucl. Phys.*, vol. 22, pp. 579–588, 1961.
- [33] Steven Weinberg, “A Model of Leptons”, *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 1967.
- [34] Abdus Salam, “Weak and Electromagnetic Interactions”, *Conf. Proc. C*, vol. 680519, pp. 367–377, 1968.
- [35] M. Bohm, H. Spiesberger, and W. Hollik, “On the One Loop Renormalization of the Electroweak Standard Model and Its Application to Leptonic Processes”, *Fortsch. Phys.*, vol. 34, pp. 687–751, 1986.
- [36] Ansgar Denner, “Techniques for calculation of electroweak radiative corrections at the one loop level and results for W physics at LEP-200”, *Fortsch. Phys.*, vol. 41, pp. 307–420, 1993.
- [37] W. F. L. Hollik, “Radiative Corrections in the Standard Model and their Role for Precision Tests of the Electroweak Theory”, *Fortsch. Phys.*, vol. 38, pp. 165–260, 1990.
- [38] M. Bohm, Ansgar Denner, and H. Joos, *Gauge theories of the strong and electroweak interaction*, 2001.
- [39] John F. Gunion, Howard E. Haber, Gordon L. Kane, and Sally Dawson, *The Higgs Hunter’s Guide*, vol. 80, 2000.
- [40] B. Grzadkowski, M. Iskrzynski, M. Misiak, and J. Rosiek, “Dimension-Six Terms in the Standard Model Lagrangian”, *JHEP*, vol. 10, pp. 085, 2010.
- [41] Nicola Cabibbo, “Unitary Symmetry and Leptonic Decays”, *Phys. Rev. Lett.*, vol. 10, pp. 531–533, 1963.
- [42] Makoto Kobayashi and Toshihide Maskawa, “CP Violation in the Renormalizable Theory of Weak Interaction”, *Prog. Theor. Phys.*, vol. 49, pp. 652–657, 1973.
- [43] Bryce S. DeWitt, “Quantum theory of gravity. ii. the manifestly covariant theory”, *Phys. Rev.*, vol. 162, pp. 1195–1239, Oct 1967.
- [44] Bryce S. DeWitt, “A GAUGE INVARIANT EFFECTIVE ACTION”, in *Oxford Conference on Quantum Gravity*, 7 1980.
- [45] Gerard ’t Hooft, “The Background Field Method in Gauge Field Theories”, in *12th Annual Winter School of Theoretical Physics*, 1975.
- [46] David G. Boulware, “Gauge dependence of the effective action”, *Phys. Rev. D*, vol. 23, pp. 389–396, Jan 1981.
- [47] L. F. Abbott, “The Background Field Method Beyond One Loop”, *Nucl. Phys. B*, vol. 185, pp. 189–203, 1981.
- [48] Ansgar Denner, Georg Weiglein, and Stefan Dittmaier, “Application of the background field method to the electroweak standard model”, *Nucl. Phys. B*, vol. 440, pp. 95–128, 1995.

- [49] Ansgar Denner and Stefan Dittmaier, “Dyson summation without violating Ward identities and the Goldstone boson equivalence theorem”, *Phys. Rev. D*, vol. 54, pp. 4499–4514, 1996.
- [50] John N. Bahcall, “Solar neutrinos. I: Theoretical”, *Phys. Rev. Lett.*, vol. 12, pp. 300–302, 1964.
- [51] R. Davis, “Solar neutrinos. II: Experimental”, *Phys. Rev. Lett.*, vol. 12, pp. 303–305, 1964.
- [52] Q. R. Ahmad et al., “Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory”, *Phys. Rev. Lett.*, vol. 87, pp. 071301, 2001.
- [53] Q. R. Ahmad et al., “Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory”, *Phys. Rev. Lett.*, vol. 89, pp. 011301, 2002.
- [54] Y. Fukuda et al., “Evidence for oscillation of atmospheric neutrinos”, *Phys. Rev. Lett.*, vol. 81, pp. 1562–1567, 1998.
- [55] R. N. Mohapatra and A. Y. Smirnov, “Neutrino Mass and New Physics”, *Ann. Rev. Nucl. Part. Sci.*, vol. 56, pp. 569–628, 2006.
- [56] Alessandro Strumia and Francesco Vissani, “Neutrino masses and mixings and...”, 6 2006.
- [57] Jack Dragos, Thomas Luu, Andrea Shindler, Jordy de Vries, and Ahmed Yousif, “Confirming the Existence of the strong CP Problem in Lattice QCD with the Gradient Flow”, *Phys. Rev. C*, vol. 103, no. 1, pp. 015202, 2021.
- [58] J. M. Pendlebury et al., “Revised experimental upper limit on the electric dipole moment of the neutron”, *Phys. Rev. D*, vol. 92, no. 9, pp. 092003, 2015.
- [59] B. Graner, Y. Chen, E.G. Lindahl, and B.R. Heckel, “Reduced limit on the permanent electric dipole moment of ^{199}Hg ”, *Physical Review Letters*, vol. 116, no. 16, April 2016.
- [60] Steven Weinberg, *The quantum theory of fields. Vol. 2: Modern applications*, Cambridge University Press, 8 2013.
- [61] F. Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln”, *Helv. Phys. Acta*, vol. 6, pp. 110–127, 1933.
- [62] Vera C. Rubin and W. Kent Ford, Jr., “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions”, *The Astrophysical Journal*, vol. 159, pp. 379, February 1970.
- [63] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov, “On the Anomalous Electroweak Baryon Number Nonconservation in the Early Universe”, *Phys. Lett. B*, vol. 155, pp. 36, 1985.
- [64] A. D. Sakharov, “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe”, *Pisma Zh. Eksp. Teor. Fiz.*, vol. 5, pp. 32–35, 1967.
- [65] W. Buchmuller and D. Wyler, “Effective Lagrangian Analysis of New Interactions and Flavor Conservation”, *Nucl. Phys. B*, vol. 268, pp. 621–653, 1986.
- [66] G. F. Giudice, C. Grojean, A. Pomarol, and R. Rattazzi, “The Strongly-Interacting Light Higgs”, *JHEP*, vol. 06, pp. 045, 2007.

- [67] F. Bonnet, M. B. Gavela, T. Ota, and W. Winter, “Anomalous Higgs couplings at the LHC, and their theoretical interpretation”, *Phys. Rev. D*, vol. 85, pp. 035016, 2012.
- [68] Tyler Corbett, O. J. P. Eboli, J. Gonzalez-Fraile, and M. C. Gonzalez-Garcia, “Constraining anomalous Higgs interactions”, *Phys. Rev. D*, vol. 86, pp. 075013, 2012.
- [69] Florian Bonnet, Toshihiko Ota, Michael Rauch, and Walter Winter, “Interpretation of precision tests in the Higgs sector in terms of physics beyond the Standard Model”, *Phys. Rev. D*, vol. 86, pp. 093014, 2012.
- [70] Giampiero Passarino, “NLO Inspired Effective Lagrangians for Higgs Physics”, *Nucl. Phys. B*, vol. 868, pp. 416–458, 2013.
- [71] Tyler Corbett, O. J. P. Eboli, J. Gonzalez-Fraile, and M. C. Gonzalez-Garcia, “Robust Determination of the Higgs Couplings: Power to the Data”, *Phys. Rev. D*, vol. 87, pp. 015022, 2013.
- [72] Roberto Contino, Margherita Ghezzi, Christophe Grojean, Margarete Muhlleitner, and Michael Spira, “Effective Lagrangian for a light Higgs-like scalar”, *JHEP*, vol. 07, pp. 035, 2013.
- [73] Stefan Dittmaier, Timo Schmidt, and Jan Schwarz, “Mixed NNLO QCD×electroweak corrections of $\mathcal{O}(N_f\alpha_s\alpha)$ to single-W/Z production at the LHC”, *JHEP*, vol. 12, pp. 201, 2020.
- [74] Ayres Freitas, W. Hollik, W. Walter, and Georg Weiglein, “Electroweak two loop corrections to the $M_W - M_Z$ mass correlation in the standard model”, *Nucl. Phys. B*, vol. 632, pp. 189–218, 2002, [Erratum: *Nucl.Phys.B* 666, 305–307 (2003)].
- [75] M. Awramik, M. Czakon, A. Onishchenko, and O. Veretin, “Bosonic corrections to Delta r at the two loop level”, *Phys. Rev. D*, vol. 68, pp. 053004, 2003.
- [76] Stefano Actis, Andrea Ferroglia, Massimo Passera, and Giampiero Passarino, “Two-Loop Renormalization in the Standard Model. Part I: Prolegomena”, *Nucl. Phys. B*, vol. 777, pp. 1–34, 2007.
- [77] Stefano Actis and Giampiero Passarino, “Two-Loop Renormalization in the Standard Model Part II: Renormalization Procedures and Computational Techniques”, *Nucl. Phys. B*, vol. 777, pp. 35–99, 2007.
- [78] Stefano Actis and Giampiero Passarino, “Two-Loop Renormalization in the Standard Model Part III: Renormalization Equations and their Solutions”, *Nucl. Phys. B*, vol. 777, pp. 100–156, 2007.
- [79] Giuseppe Degrandi and Alessandro Vicini, “Two loop renormalization of the electric charge in the standard model”, *Phys. Rev.*, vol. D69, pp. 073007, 2004.
- [80] P. Breitenlohner and D. Maison, “Dimensional Renormalization and the Action Principle”, *Commun. Math. Phys.*, vol. 52, pp. 11–38, 1977.
- [81] P. Breitenlohner and D. Maison, “Dimensionally Renormalized Green’s Functions for Theories with Massless Particles. 1.”, *Commun. Math. Phys.*, vol. 52, pp. 39, 1977.
- [82] P. Breitenlohner and D. Maison, “Dimensionally Renormalized Green’s Functions for Theories with Massless Particles. 2.”, *Commun. Math. Phys.*, vol. 52, pp. 55, 1977.
- [83] Gerard ’t Hooft, “Dimensional regularization and the renormalization group”, *Nucl. Phys. B*, vol. 61, pp. 455–468, 1973.

- [84] Dirk Kreimer, “The $\gamma(5)$ Problem and Anomalies: A Clifford Algebra Approach”, *Phys. Lett. B*, vol. 237, pp. 59–62, 1990.
- [85] Gerard 't Hooft, “Renormalization of Massless Yang-Mills Fields”, *Nucl. Phys. B*, vol. 33, pp. 173–199, 1971.
- [86] Gerard 't Hooft, “Renormalizable Lagrangians for Massive Yang-Mills Fields”, *Nucl. Phys. B*, vol. 35, pp. 167–188, 1971.
- [87] Gerard 't Hooft and M. J. G. Veltman, “Regularization and Renormalization of Gauge Fields”, *Nucl. Phys. B*, vol. 44, pp. 189–213, 1972.
- [88] B. W. Lee and Jean Zinn-Justin, “Spontaneously Broken Gauge Symmetries Part 2: Perturbation Theory and Renormalization”, *Phys. Rev. D*, vol. 5, pp. 3137–3155, 1972, [Erratum: *Phys. Rev. D* 8, 4654 (1973)].
- [89] Benjamin W. Lee, “Renormalization of Gauge Theories: Unbroken and Broken”, *Phys. Rev. D*, vol. 9, pp. 933–946, 1974.
- [90] D. A. Ross and J. C. Taylor, “Renormalization of a unified theory of weak and electromagnetic interactions”, *Nucl. Phys. B*, vol. 51, pp. 125–144, 1973, [Erratum: *Nucl. Phys. B* 58, 643–643 (1973)].
- [91] A. Sirlin, “Radiative Corrections in the $SU(2)_L \times U(1)$ Theory: A Simple Renormalization Framework”, *Phys. Rev.*, vol. D22, pp. 971–981, 1980.
- [92] K. I. Aoki, Z. Hioki, M. Konuma, R. Kawabe, and T. Muta, “Electroweak Theory. Framework of On-Shell Renormalization and Study of Higher Order Effects”, *Prog. Theor. Phys. Suppl.*, vol. 73, pp. 1–225, 1982.
- [93] W. J. Marciano and A. Sirlin, “Radiative Corrections to Neutrino Induced Neutral Current Phenomena in the $SU(2)_L \times U(1)$ Theory”, *Phys. Rev. D*, vol. 22, pp. 2695, 1980, [Erratum: *Phys. Rev. D* 31, 213 (1985)].
- [94] Stefan Dittmaier, “Electric charge renormalization to all orders”, *Phys. Rev. D*, vol. 103, no. 5, pp. 053006, 2021.
- [95] Fred Jegerlehner, “Hadronic contributions to the photon vacuum polarization and their role in precision physics”, *J. Phys. G*, vol. 29, pp. 101–110, 2003.
- [96] A. Sirlin and R. Zucchini, “Dependence of the Quartic Coupling $H(m)$ on $M(H)$ and the Possible Onset of New Physics in the Higgs Sector of the Standard Model”, *Nucl. Phys. B*, vol. 266, pp. 389–409, 1986.
- [97] Ansgar Denner, S. Dittmaier, M. Roth, and D. Wackeroth, “Predictions for all processes $e^+ e^- \rightarrow 4$ fermions + gamma”, *Nucl. Phys. B*, vol. 560, pp. 33–65, 1999.
- [98] Ansgar Denner, S. Dittmaier, M. Roth, and L. H. Wieders, “Electroweak corrections to charged-current $e^+ e^- \rightarrow 4$ fermion processes: Technical details and further results”, *Nucl. Phys. B*, vol. 724, pp. 247–294, 2005, [Erratum: *Nucl. Phys. B* 854, 504–507 (2012)].

- [99] Ansgar Denner and S. Dittmaier, “The Complex-mass scheme for perturbative calculations with unstable particles”, *Nucl. Phys. B Proc. Suppl.*, vol. 160, pp. 22–26, 2006.
- [100] John C. Collins, *Renormalization : An Introduction to Renormalization, the Renormalization Group and the Operator-Product Expansion*, vol. 26 of *Cambridge Monographs on Mathematical Physics*, Cambridge University Press, Cambridge, 1984.
- [101] J. G. Korner, D. Kreimer, and K. Schilcher, “A Practicable $\gamma(5)$ scheme in dimensional regularization”, *Z. Phys. C*, vol. 54, pp. 503–512, 1992.
- [102] Z. Bern and A. G. Morgan, “Massive loop amplitudes from unitarity”, *Nucl. Phys. B*, vol. 467, pp. 479–509, 1996.
- [103] Zvi Bern, Abilio De Freitas, and Lance J. Dixon, “Two loop helicity amplitudes for gluon-gluon scattering in QCD and supersymmetric Yang-Mills theory”, *JHEP*, vol. 03, pp. 018, 2002.
- [104] Radja Boughezal, Frank Petriello, Ulrich Schubert, and Hongxi Xing, “Spin-dependent quark beam function at NNLO”, *Phys. Rev. D*, vol. 96, no. 3, pp. 034001, 2017.
- [105] Vladimir A. Smirnov, “Evaluating Feynman integrals”, *Springer Tracts Mod. Phys.*, vol. 211, pp. 1–244, 2004.
- [106] Martijn Hidding, “DiffExp, a Mathematica package for computing Feynman integrals in terms of one-dimensional series expansions”, *Comput. Phys. Commun.*, vol. 269, pp. 108125, 2021.
- [107] S. Laporta, “High-precision calculation of multiloop Feynman integrals by difference equations”, *Int. J. Mod. Phys. A*, vol. 15, pp. 5087–5159, 2000.
- [108] Charalampos Anastasiou and Achilleas Lazopoulos, “Automatic integral reduction for higher order perturbative calculations”, *JHEP*, vol. 07, pp. 046, 2004.
- [109] Xin Guan, Xiao Liu, Yan-Qing Ma, and Wen-Hao Wu, “Blade: A package for block-triangular form improved Feynman integrals decomposition”, 5 2024.
- [110] A. V. Smirnov, “Algorithm FIRE – Feynman Integral REduction”, *JHEP*, vol. 10, pp. 107, 2008.
- [111] Alexander V. Smirnov and Mao Zeng, “FIRE 6.5: Feynman integral reduction with new simplification library”, *Comput. Phys. Commun.*, vol. 302, pp. 109261, 2024.
- [112] Philipp Maierhöfer, Johann Usovitsch, and Peter Uwer, “Kira—A Feynman integral reduction program”, *Comput. Phys. Commun.*, vol. 230, pp. 99–112, 2018.
- [113] Jonas Klappert, Fabian Lange, Philipp Maierhöfer, and Johann Usovitsch, “Integral reduction with Kira 2.0 and finite field methods”, *Comput. Phys. Commun.*, vol. 266, pp. 108024, 2021.
- [114] Fabian Lange, Johann Usovitsch, and Zihao Wu, “Kira 3: integral reduction with efficient seeding and optimized equation selection”, 5 2025.
- [115] Roman N. Lee, “LiteRed 1.4: a powerful tool for reduction of multiloop integrals”, *J. Phys. Conf. Ser.*, vol. 523, pp. 012059, 2014.

- [116] Zihao Wu, Janko Boehm, Rourou Ma, Hefeng Xu, and Yang Zhang, “NeatIBP 1.0, a package generating small-size integration-by-parts relations for Feynman integrals”, *Comput. Phys. Commun.*, vol. 295, pp. 108999, 2024.
- [117] A. von Manteuffel and C. Studerus, “Reduze 2 - Distributed Feynman Integral Reduction”, 1 2012.
- [118] Andreas von Manteuffel and Robert M. Schabinger, “A novel approach to integration by parts reduction”, *Phys. Lett. B*, vol. 744, pp. 101–104, 2015.
- [119] Tiziano Peraro, “FiniteFlow: multivariate functional reconstruction using finite fields and dataflow graphs”, *JHEP*, vol. 07, pp. 031, 2019.
- [120] Jonas Klappert and Fabian Lange, “Reconstructing rational functions with FireFly”, *Comput. Phys. Commun.*, vol. 247, pp. 106951, 2020.
- [121] Vitaly Magerya, “Rational Tracer: a Tool for Faster Rational Function Reconstruction”, 11 2022.
- [122] P. Mastrolia and E. Remiddi, “Analytic evaluation of Feynman graph integrals”, *Nucl. Phys. B Proc. Suppl.*, vol. 116, pp. 412–416, 2003.
- [123] M. Czakor and H. Czyz, “Why and how to use a differential equation method to calculate multiloop integrals”, *Acta Phys. Polon. B*, vol. 32, pp. 3823, 2001.
- [124] Michele Caffo, H. Czyz, S. Laporta, and E. Remiddi, “The Master differential equations for the two loop sunrise selfmass amplitudes”, *Nuovo Cim. A*, vol. 111, pp. 365–389, 1998.
- [125] A. V. Kotikov, “The Differential equation method: Evaluation of complicated Feynman diagrams”, in *15th International Workshop on High-Energy Physics and Quantum Field Theory (QFTHEP 2000)*, 7 2000, pp. 203–210.
- [126] V. A. Smirnov, *Renormalization and asymptotic expansions*, vol. 14, 1991.
- [127] M. Beneke and Vladimir A. Smirnov, “Asymptotic expansion of Feynman integrals near threshold”, *Nucl. Phys. B*, vol. 522, pp. 321–344, 1998.
- [128] G. Heinrich, S. Jahn, S. P. Jones, M. Kerner, F. Langer, V. Magerya, A. Pöldaru, J. Schlenk, and E. Villa, “Expansion by regions with pySecDec”, *Comput. Phys. Commun.*, vol. 273, pp. 108267, 2022.
- [129] S. Borowka, G. Heinrich, S. Jahn, S. P. Jones, M. Kerner, J. Schlenk, and T. Zirke, “pySecDec: a toolbox for the numerical evaluation of multi-scale integrals”, *Comput. Phys. Commun.*, vol. 222, pp. 313–326, 2018.
- [130] Xiao Liu and Yan-Qing Ma, “AMFlow: A Mathematica package for Feynman integrals computation via auxiliary mass flow”, *Comput. Phys. Commun.*, vol. 283, pp. 108565, 2023.
- [131] Johannes M. Henn, “Multiloop integrals in dimensional regularization made simple”, *Phys. Rev. Lett.*, vol. 110, pp. 251601, 2013.
- [132] Francesco Moriello, “Generalised power series expansions for the elliptic planar families of Higgs + jet production at two loops”, *JHEP*, vol. 01, pp. 150, 2020.

- [133] U. Aglietti and R. Bonciani, “Master integrals with one massive propagator for the two loop electroweak form-factor”, *Nucl. Phys. B*, vol. 668, pp. 3–76, 2003.
- [134] U. Aglietti and R. Bonciani, “Master integrals with 2 and 3 massive propagators for the 2 loop electroweak form-factor - planar case”, *Nucl. Phys. B*, vol. 698, pp. 277–318, 2004.
- [135] T. Gehrmann, T. Huber, and D. Maitre, “Two-loop quark and gluon form-factors in dimensional regularisation”, *Phys. Lett. B*, vol. 622, pp. 295–302, 2005.
- [136] R. Bonciani, A. Ferroglia, T. Gehrmann, D. Maitre, and C. Studerus, “Two-Loop Fermionic Corrections to Heavy-Quark Pair Production: The Quark-Antiquark Channel”, *JHEP*, vol. 07, pp. 129, 2008.
- [137] R. Bonciani, A. Ferroglia, T. Gehrmann, and C. Studerus, “Two-Loop Planar Corrections to Heavy-Quark Pair Production in the Quark-Antiquark Channel”, *JHEP*, vol. 08, pp. 067, 2009.
- [138] Roberto Bonciani, Stefano Di Vita, Pierpaolo Mastrolia, and Ulrich Schubert, “Two-Loop Master Integrals for the mixed EW-QCD virtual corrections to Drell-Yan scattering”, *JHEP*, vol. 09, pp. 091, 2016.
- [139] Matthias Heller, Andreas von Manteuffel, and Robert M. Schabinger, “Multiple polylogarithms with algebraic arguments and the two-loop EW-QCD Drell-Yan master integrals”, *Phys. Rev. D*, vol. 102, no. 1, pp. 016025, 2020.
- [140] Syed Mehedi Hasan and Ulrich Schubert, “Master Integrals for the mixed QCD-QED corrections to the Drell-Yan production of a massive lepton pair”, *JHEP*, vol. 11, pp. 107, 2020.
- [141] Kuo-Tsai Chen, “Iterated path integrals”, *Bull. Am. Math. Soc.*, vol. 83, pp. 831–879, 1977.
- [142] L. Naterop, A. Signer, and Y. Ulrich, “handyG —Rapid numerical evaluation of generalised polylogarithms in Fortran”, *Comput. Phys. Commun.*, vol. 253, pp. 107165, 2020.
- [143] Jakob Ablinger, “A Computer Algebra Toolbox for Harmonic Sums Related to Particle Physics”, *Master’s thesis, Linz U.*, 2009.
- [144] Jakob Ablinger, “The package HarmonicSums: Computer Algebra and Analytic aspects of Nested Sums”, *PoS*, vol. LL2014, pp. 019, 2014.
- [145] Claude Duhr and Falko Dulat, “PolyLogTools — polylogs for the masses”, *JHEP*, vol. 08, pp. 135, 2019.
- [146] T. Hahn and M. Perez-Victoria, “Automatized one loop calculations in four-dimensions and D-dimensions”, *Comput. Phys. Commun.*, vol. 118, pp. 153–165, 1999.
- [147] Johannes Blumlein, “The Theory of Deeply Inelastic Scattering”, *Prog. Part. Nucl. Phys.*, vol. 69, pp. 28–84, 2013.
- [148] P. Jimenez-Delgado, W. Melnitchouk, and J. F. Owens, “Parton momentum and helicity distributions in the nucleon”, *J. Phys. G*, vol. 40, pp. 093102, 2013.
- [149] Stefano Forte and Graeme Watt, “Progress in the Determination of the Partonic Structure of the Proton”, *Ann. Rev. Nucl. Part. Sci.*, vol. 63, pp. 291–328, 2013.

- [150] Jun Gao, Lucian Harland-Lang, and Juan Rojo, “The Structure of the Proton in the LHC Precision Era”, *Phys. Rept.*, vol. 742, pp. 1–121, 2018.
- [151] Jacob J. Ethier and Emanuele R. Nocera, “Parton Distributions in Nucleons and Nuclei”, *Ann. Rev. Nucl. Part. Sci.*, vol. 70, pp. 43–76, 2020.
- [152] Guido Altarelli and G. Parisi, “Asymptotic Freedom in Parton Language”, *Nucl. Phys. B*, vol. 126, pp. 298–318, 1977.
- [153] Yu. L. Dokshitser, “Calculation of structure functions of deep-inelastic scattering and $e+e-$ annihilation by perturbation theory in quantum chromodynamics”, *Sov. Phys. JETP*, vol. 46, pp. 641–653, 1977.
- [154] S. Frixione, Z. Kunszt, and A. Signer, “Three jet cross-sections to next-to-leading order”, *Nucl. Phys. B*, vol. 467, pp. 399–442, 1996.
- [155] S. Catani and M. H. Seymour, “The Dipole formalism for the calculation of QCD jet cross-sections at next-to-leading order”, *Phys. Lett. B*, vol. 378, pp. 287–301, 1996.
- [156] S. Catani and M. H. Seymour, “A General algorithm for calculating jet cross-sections in NLO QCD”, *Nucl. Phys. B*, vol. 485, pp. 291–419, 1997, [Erratum: *Nucl.Phys.B* 510, 503–504 (1998)].
- [157] Stefano Catani and Massimiliano Grazzini, “An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC”, *Phys. Rev. Lett.*, vol. 98, pp. 222002, 2007.
- [158] Stefano Camarda, Leandro Cieri, and Giancarlo Ferrera, “Fiducial perturbative power corrections within the q_T subtraction formalism”, *Eur. Phys. J. C*, vol. 82, no. 6, pp. 575, 2022.
- [159] Luca Buonocore, Stefan Kallweit, Luca Rottoli, and Marius Wiesemann, “Linear power corrections for two-body kinematics in the q_T subtraction formalism”, *Phys. Lett. B*, vol. 829, pp. 137118, 2022.
- [160] R. Frederix, S. Frixione, V. Hirschi, D. Pagani, H. S. Shao, and M. Zaro, “The automation of next-to-leading order electroweak calculations”, *JHEP*, vol. 07, pp. 185, 2018.
- [161] Thomas Hahn, “Generating Feynman diagrams and amplitudes with FeynArts 3”, *Comput. Phys. Commun.*, vol. 140, pp. 418–431, 2001.
- [162] Massimiliano Grazzini, Stefan Kallweit, and Marius Wiesemann, “Fully differential NNLO computations with MATRIX”, *Eur. Phys. J. C*, vol. 78, no. 7, pp. 537, 2018.
- [163] Federico Buccioni, Jean-Nicolas Lang, Jonas M. Lindert, Philipp Maierhöfer, Stefano Pozzorini, Hantian Zhang, and Max F. Zoller, “OpenLoops 2”, *Eur. Phys. J. C*, vol. 79, no. 10, pp. 866, 2019.
- [164] Stefano Actis, Ansgar Denner, Lars Hofer, Jean-Nicolas Lang, Andreas Scharf, and Sandro Uccirati, “RECOLA: REcursive Computation of One-Loop Amplitudes”, *Comput. Phys. Commun.*, vol. 214, pp. 140–173, 2017.
- [165] S.D. Drell and Tung-Mow Yan, “Massive Lepton Pair Production in Hadron-Hadron Collisions at High-Energies”, *Phys. Rev. Lett.*, vol. 25, pp. 316–320, 1970, [Erratum: *Phys.Rev.Lett.* 25, 902 (1970)].

- [166] Tevatron Electroweak Working Group, “2012 Update of the Combination of CDF and D0 Results for the Mass of the W Boson”, 4 2012.
- [167] Morad Aaboud et al., “Measurement of the W -boson mass in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”, *Eur. Phys. J. C*, vol. 78, no. 2, pp. 110, 2018, [Erratum: *Eur.Phys.J.C* 78, 898 (2018)].
- [168] Timo Antero Aaltonen et al., “Tevatron Run II combination of the effective leptonic electroweak mixing angle”, *Phys. Rev. D*, vol. 97, no. 11, pp. 112007, 2018.
- [169] “Measurement of the effective leptonic weak mixing angle using electron and muon pairs from Z -boson decay in the ATLAS experiment at $\sqrt{s} = 8$ TeV”, *ATLAS-CONF-2018-037*, 7 2018.
- [170] Aram Hayrapetyan et al., “Measurement of the Drell–Yan forward-backward asymmetry and of the effective leptonic weak mixing angle in proton-proton collisions at $s=13\text{TeV}$ ”, *Phys. Lett. B*, vol. 866, pp. 139526, 2025.
- [171] Carlo Michel Carloni Calame, Mauro Chiesa, Homero Martinez, Guido Montagna, Oreste Nicosini, Fulvio Piccinini, and Alessandro Vicini, “Precision Measurement of the W-Boson Mass: Theoretical Contributions and Uncertainties”, *Phys. Rev. D*, vol. 96, no. 9, pp. 093005, 2017.
- [172] E. Bagnaschi and A. Vicini, “Parton Density Uncertainties and the Determination of Electroweak Parameters at Hadron Colliders”, *Phys. Rev. Lett.*, vol. 126, no. 4, pp. 041801, 2021.
- [173] Arnd Behring, Federico Buccioni, Fabrizio Caola, Maximilian Delto, Matthieu Jaquier, Kirill Melnikov, and Raoul Röntsch, “Estimating the impact of mixed QCD-electroweak corrections on the W -mass determination at the LHC”, *Phys. Rev. D*, vol. 103, no. 11, pp. 113002, 2021.
- [174] Luca Rottoli, Paolo Torrielli, and Alessandro Vicini, “Determination of the W-boson mass at hadron colliders”, *Eur. Phys. J. C*, vol. 83, no. 10, pp. 948, 2023.
- [175] Paolo Torrielli, Luca Rottoli, and Alessandro Vicini, “A new observable for W-mass determination”, *PoS*, vol. RADCOR2023, pp. 038, 2024.
- [176] S. Amoroso et al., “Snowmass 2021 Whitepaper: Proton Structure at the Precision Frontier”, *Acta Phys. Polon. B*, vol. 53, no. 12, pp. 12–A1, 2022.
- [177] Richard D. Ball, Stefano Carrazza, Luigi Del Debbio, Stefano Forte, Zahari Kassabov, Juan Rojo, Emma Slade, and Maria Ubiali, “Precision determination of the strong coupling constant within a global PDF analysis”, *Eur. Phys. J. C*, vol. 78, no. 5, pp. 408, 2018.
- [178] Georges Aad et al., “A precise determination of the strong-coupling constant from the recoil of Z bosons with the ATLAS experiment at $\sqrt{s} = 8$ TeV”, 9 2023.
- [179] Guido Altarelli, R.Keith Ellis, and G. Martinelli, “Large Perturbative Corrections to the Drell-Yan Process in QCD”, *Nucl. Phys. B*, vol. 157, pp. 461–497, 1979.
- [180] R. Hamberg, W.L. van Neerven, and T. Matsuura, “A complete calculation of the order α_s^2 correction to the Drell-Yan K factor”, *Nucl. Phys. B*, vol. 359, pp. 343–405, 1991, [Erratum: *Nucl.Phys.B* 644, 403–404 (2002)].

- [181] Robert V. Harlander and William B. Kilgore, “Next-to-next-to-leading order Higgs production at hadron colliders”, *Phys. Rev. Lett.*, vol. 88, pp. 201801, 2002.
- [182] Charalampos Anastasiou, Lance J. Dixon, Kirill Melnikov, and Frank Petriello, “Dilepton rapidity distribution in the Drell-Yan process at NNLO in QCD”, *Phys. Rev. Lett.*, vol. 91, pp. 182002, 2003.
- [183] Kirill Melnikov and Frank Petriello, “Electroweak gauge boson production at hadron colliders through $\mathcal{O}(\alpha_s^2)$ ”, *Phys. Rev. D*, vol. 74, pp. 114017, 2006.
- [184] Stefano Catani, Leandro Cieri, Giancarlo Ferrera, Daniel de Florian, and Massimiliano Grazzini, “Vector boson production at hadron colliders: a fully exclusive QCD calculation at NNLO”, *Phys. Rev. Lett.*, vol. 103, pp. 082001, 2009.
- [185] Stefano Catani, Giancarlo Ferrera, and Massimiliano Grazzini, “W Boson Production at Hadron Colliders: The Lepton Charge Asymmetry in NNLO QCD”, *JHEP*, vol. 05, pp. 006, 2010.
- [186] Stefan Dittmaier and Michael Krämer, “Electroweak radiative corrections to W boson production at hadron colliders”, *Phys. Rev. D*, vol. 65, pp. 073007, 2002.
- [187] U. Baur and D. Wackeroth, “Electroweak radiative corrections to $p\bar{p} \rightarrow W^\pm \rightarrow \ell^\pm \nu$ beyond the pole approximation”, *Phys. Rev. D*, vol. 70, pp. 073015, 2004.
- [188] V.A. Zykunov, “Radiative corrections to the Drell-Yan process at large dilepton invariant masses”, *Phys. Atom. Nucl.*, vol. 69, pp. 1522, 2006.
- [189] A. Arbuzov, D. Bardin, S. Bondarenko, P. Christova, L. Kalinovskaya, G. Nanava, and R. Sadykov, “One-loop corrections to the Drell-Yan process in SANC. I. The Charged current case”, *Eur. Phys. J. C*, vol. 46, pp. 407–412, 2006, [Erratum: *Eur.Phys.J.C* 50, 505 (2007)].
- [190] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and A. Vicini, “Precision electroweak calculation of the charged current Drell-Yan process”, *JHEP*, vol. 12, pp. 016, 2006.
- [191] U. Baur, O. Brein, W. Hollik, C. Schappacher, and D. Wackeroth, “Electroweak radiative corrections to neutral current Drell-Yan processes at hadron colliders”, *Phys. Rev. D*, vol. 65, pp. 033007, 2002.
- [192] V.A. Zykunov, “Weak radiative corrections to Drell-Yan process for large invariant mass of dilepton pair”, *Phys. Rev. D*, vol. 75, pp. 073019, 2007.
- [193] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and A. Vicini, “Precision electroweak calculation of the production of a high transverse-momentum lepton pair at hadron colliders”, *JHEP*, vol. 10, pp. 109, 2007.
- [194] A. Arbuzov, D. Bardin, S. Bondarenko, P. Christova, L. Kalinovskaya, G. Nanava, and R. Sadykov, “One-loop corrections to the Drell-Yan process in SANC. (II). The Neutral current case”, *Eur. Phys. J. C*, vol. 54, pp. 451–460, 2008.
- [195] Stefan Dittmaier and Max Huber, “Radiative corrections to the neutral-current Drell-Yan process in the Standard Model and its minimal supersymmetric extension”, *JHEP*, vol. 01, pp. 060, 2010.

- [196] Claude Duhr, Falko Dulat, and Bernhard Mistlberger, “Drell-Yan Cross Section to Third Order in the Strong Coupling Constant”, *Phys. Rev. Lett.*, vol. 125, no. 17, pp. 172001, 2020.
- [197] Xuan Chen, Thomas Gehrmann, Nigel Glover, Alexander Huss, Tong-Zhi Yang, and Hua Xing Zhu, “Di-lepton Rapidity Distribution in Drell-Yan Production to Third Order in QCD”, 7 2021.
- [198] Claude Duhr, Falko Dulat, and Bernhard Mistlberger, “Charged current Drell-Yan production at N^3LO ”, *JHEP*, vol. 11, pp. 143, 2020.
- [199] Stefano Camarda, Leandro Cieri, and Giancarlo Ferrera, “Drell-Yan lepton-pair production: q_T resummation at N^3LL accuracy and fiducial cross sections at N^3LO ”, 3 2021.
- [200] Xuan Chen, Thomas Gehrmann, E. W. N. Glover, Alexander Huss, Pier Francesco Monni, Emanuele Re, Luca Rottoli, and Paolo Torrielli, “Third-Order Fiducial Predictions for Drell-Yan Production at the LHC”, *Phys. Rev. Lett.*, vol. 128, no. 25, pp. 252001, 2022.
- [201] Tobias Neumann and John Campbell, “Fiducial Drell-Yan production at the LHC improved by transverse-momentum resummation at $N4LLp+N3LO$ ”, *Phys. Rev. D*, vol. 107, no. 1, pp. L011506, 2023.
- [202] John Campbell and Tobias Neumann, “Third order QCD predictions for fiducial W-boson production”, *JHEP*, vol. 11, pp. 127, 2023.
- [203] Daniel de Florian, Manuel Der, and Ignacio Fabre, “ $QCD \oplus QED$ NNLO corrections to Drell Yan production”, *Phys. Rev. D*, vol. 98, no. 9, pp. 094008, 2018.
- [204] Leandro Cieri, Daniel de Florian, Manuel Der, and Javier Mazzitelli, “Mixed $QCD \otimes QED$ corrections to exclusive Drell Yan production using the q_T -subtraction method”, *JHEP*, vol. 09, pp. 155, 2020.
- [205] Maximilian Delto, Matthieu Jaquier, Kirill Melnikov, and Raoul Röntschi, “Mixed $QCD \otimes QED$ corrections to on-shell Z boson production at the LHC”, *JHEP*, vol. 01, pp. 043, 2020.
- [206] Roberto Bonciani, Federico Buccioni, Roberto Mondini, and Alessandro Vicini, “Double-real corrections at $\mathcal{O}(\alpha\alpha_s)$ to single gauge boson production”, *Eur. Phys. J. C*, vol. 77, no. 3, pp. 187, 2017.
- [207] Roberto Bonciani, Federico Buccioni, Narayan Rana, Ilario Triscari, and Alessandro Vicini, “NNLO $QCD \times EW$ corrections to Z production in the $q\bar{q}$ channel”, *Phys. Rev. D*, vol. 101, no. 3, pp. 031301, 2020.
- [208] Roberto Bonciani, Federico Buccioni, Narayan Rana, and Alessandro Vicini, “NNLO $QCD \times EW$ corrections to on-shell Z production”, *Phys. Rev. Lett.*, vol. 125, no. 23, pp. 232004, 2020.
- [209] Federico Buccioni, Fabrizio Caola, Maximilian Delto, Matthieu Jaquier, Kirill Melnikov, and Raoul Röntschi, “Mixed QCD -electroweak corrections to on-shell Z production at the LHC”, *Phys. Lett. B*, vol. 811, pp. 135969, 2020.
- [210] Arnd Behring, Federico Buccioni, Fabrizio Caola, Maximilian Delto, Matthieu Jaquier, Kirill Melnikov, and Raoul Röntschi, “Mixed QCD -electroweak corrections to W -boson production in hadron collisions”, *Phys. Rev. D*, vol. 103, no. 1, pp. 013008, 2021.

- [211] Roberto Bonciani, Federico Buccioni, Narayan Rana, and Alessandro Vicini, “On-shell Z boson production at hadron colliders through $\mathcal{O}(\alpha\alpha_s)$ ”, *JHEP*, vol. 02, pp. 095, 2022.
- [212] Stefan Dittmaier, Alexander Huss, and Christian Schwinn, “Mixed QCD-electroweak $\mathcal{O}(\alpha_s\alpha)$ corrections to Drell-Yan processes in the resonance region: pole approximation and non-factorizable corrections”, *Nucl. Phys. B*, vol. 885, pp. 318–372, 2014.
- [213] Stefan Dittmaier, Alexander Huss, and Christian Schwinn, “Dominant mixed QCD-electroweak $\mathcal{O}(\alpha_s\alpha)$ corrections to Drell–Yan processes in the resonance region”, *Nucl. Phys. B*, vol. 904, pp. 216–252, 2016.
- [214] Stefan Dittmaier, Alexander Huss, and Jan Schwarz, “Mixed NNLO QCD $\times$ electroweak corrections to single-Z production in pole approximation: differential distributions and forward-backward asymmetry”, *JHEP*, vol. 05, pp. 170, 2024.
- [215] Luca Buonocore, Massimiliano Grazzini, Stefan Kallweit, Chiara Savoini, and Francesco Tramontano, “Mixed QCD-EW corrections to $pp \rightarrow \ell\nu_\ell + X$ at the LHC”, *Phys. Rev. D*, vol. 103, pp. 114012, 2021.
- [216] Roberto Bonciani, Luca Buonocore, Massimiliano Grazzini, Stefan Kallweit, Narayan Rana, Francesco Tramontano, and Alessandro Vicini, “Mixed Strong-Electroweak Corrections to the Drell-Yan Process”, *Phys. Rev. Lett.*, vol. 128, no. 1, pp. 012002, 2022.
- [217] Federico Buccioni, Fabrizio Caola, Herschel A. Chawdhry, Federica Devoto, Matthias Heller, Andreas von Manteuffel, Kirill Melnikov, Raoul Röntsch, and Chiara Signorile-Signorile, “Mixed QCD-electroweak corrections to dilepton production at the LHC in the high invariant mass region”, 3 2022.
- [218] Matthias Heller, Andreas von Manteuffel, Robert M. Schabinger, and Hubert Spiesberger, “Mixed EW-QCD two-loop amplitudes for $q\bar{q} \rightarrow \ell^+\ell^-$ and γ_5 scheme independence of multi-loop corrections”, *JHEP*, vol. 05, pp. 213, 2021.
- [219] G. Balossini, G. Montagna, C. M. Carloni Calame, M. Moretti, M. Treccani, O. Nicrosini, F. Piccinini, and A. Vicini, “Electroweak & QCD corrections to Drell Yan processes”, *Acta Phys. Polon. B*, vol. 39, pp. 1675, 2008.
- [220] Giovanni Balossini, Guido Montagna, Carlo Michel Carloni Calame, Mauro Moretti, Oreste Nicrosini, Fulvio Piccinini, Michele Treccani, and Alessandro Vicini, “Combination of electroweak and QCD corrections to single W production at the Fermilab Tevatron and the CERN LHC”, *JHEP*, vol. 01, pp. 013, 2010.
- [221] Massimiliano Grazzini, S. Kallweit, Jonas M. Lindert, Stefano Pozzorini, and Marius Wiesemann, “NNLO QCD + NLO EW with Matrix+OpenLoops: precise predictions for vector-boson pair production”, *JHEP*, vol. 02, pp. 087, 2020.
- [222] Daniel de Florian, Germán F. R. Sborlini, and Germán Rodrigo, “QED corrections to the Altarelli–Parisi splitting functions”, *Eur. Phys. J. C*, vol. 76, no. 5, pp. 282, 2016.
- [223] Daniel de Florian, Germán F. R. Sborlini, and Germán Rodrigo, “Two-loop QED corrections to the Altarelli–Parisi splitting functions”, *JHEP*, vol. 10, pp. 056, 2016.

- [224] Keping Xie, T. J. Hobbs, Tie-Jiun Hou, Carl Schmidt, Mengshi Yan, and C. P. Yuan, “Photon PDF within the CT18 global analysis”, *Phys. Rev. D*, vol. 105, no. 5, pp. 054006, 2022.
- [225] T. Cridge, L. A. Harland-Lang, and R. S. Thorne, “Combining QED and Approximate N³LO QCD Corrections in a Global PDF Fit: MSHT20qed_an3lo PDFs”, 12 2023.
- [226] Andrea Barontini, Niccolo Laurenti, and Juan Rojo, “NNPDF4.0 aN³LO PDFs with QED corrections”, in *31st International Workshop on Deep-Inelastic Scattering and Related Subjects*, 6 2024.
- [227] M. L. Mangano et al., “Physics at a 100 TeV pp Collider: Standard Model Processes”, 7 2016.
- [228] Aneesh Manohar, Paolo Nason, Gavin P. Salam, and Giulia Zanderighi, “How bright is the proton? A precise determination of the photon parton distribution function”, *Phys. Rev. Lett.*, vol. 117, no. 24, pp. 242002, 2016.
- [229] Aneesh V. Manohar, Paolo Nason, Gavin P. Salam, and Giulia Zanderighi, “The Photon Content of the Proton”, *JHEP*, vol. 12, pp. 046, 2017.
- [230] Richard D. Ball, Alessandro Candido, Stefano Forte, Felix Hekhorn, Emanuele R. Nocera, Juan Rojo, and Christopher Schwan, “Parton distributions and new physics searches: the Drell–Yan forward–backward asymmetry as a case study”, *Eur. Phys. J. C*, vol. 82, no. 12, pp. 1160, 2022.
- [231] Luca Barze, Guido Montagna, Paolo Nason, Oreste Nicrosini, Fulvio Piccinini, and Alessandro Vicini, “Neutral current Drell-Yan with combined QCD and electroweak corrections in the POWHEG BOX”, *Eur. Phys. J. C*, vol. 73, no. 6, pp. 2474, 2013.
- [232] Stefan Kallweit, Jonas M. Lindert, Philipp Maierhöfer, Stefano Pozzorini, and Marek Schönherr, “NLO QCD+EW predictions for V + jets including off-shell vector-boson decays and multijet merging”, *JHEP*, vol. 04, pp. 021, 2016.
- [233] J. R. Andersen et al., “Les Houches 2015: Physics at TeV Colliders Standard Model Working Group Report”, in *9th Les Houches Workshop on Physics at TeV Colliders*, 5 2016.
- [234] S. Alioli et al., “Precision studies of observables in $pp \rightarrow W \rightarrow l\nu_l$ and $pp \rightarrow \gamma, Z \rightarrow l^+l^-$ processes at the LHC”, *Eur. Phys. J. C*, vol. 77, no. 5, pp. 280, 2017.
- [235] Mauro Chiesa, Fulvio Piccinini, and Alessandro Vicini, “Direct determination of $\sin^2 \theta_{eff}^\ell$ at hadron colliders”, *Phys. Rev. D*, vol. 100, no. 7, pp. 071302, 2019.
- [236] Bernd A. Kniehl, “Two Loop Corrections to the Vacuum Polarizations in Perturbative QCD”, *Nucl. Phys. B*, vol. 347, pp. 86–104, 1990.
- [237] A. Djouadi and P. Gambino, “Electroweak gauge bosons selfenergies: Complete QCD corrections”, *Phys. Rev. D*, vol. 49, pp. 3499–3511, 1994, [Erratum: *Phys.Rev.D* 53, 4111 (1996)].
- [238] Stefano Catani, “The Singular behavior of QCD amplitudes at two loop order”, *Phys. Lett. B*, vol. 427, pp. 161–171, 1998.
- [239] Paulo Nogueira, “Automatic Feynman graph generation”, *J. Comput. Phys.*, vol. 105, pp. 279–289, 1993.

- [240] J. A. M. Vermaseren, “New features of FORM”, 10 2000.
- [241] Wolfram Research, Inc., “Mathematica, Version 12.3.1”, Champaign, IL, 2021.
- [242] Mario Argeri and Pierpaolo Mastrolia, “Feynman Diagrams and Differential Equations”, *Int. J. Mod. Phys. A*, vol. 22, pp. 4375–4436, 2007.
- [243] R. N. Lee, “Presenting LiteRed: a tool for the Loop InTEgrals REDuction”, 12 2012.
- [244] C. Studerus, “Reduze-Feynman Integral Reduction in C++”, *Comput. Phys. Commun.*, vol. 181, pp. 1293–1300, 2010.
- [245] J. Frenkel and J. C. Taylor, “Exponentiation of Leading Infrared Divergences in Massless Yang-Mills Theories”, *Nucl. Phys. B*, vol. 116, pp. 185–194, 1976.
- [246] Stefano Catani, Stefan Dittmaier, Michael H. Seymour, and Zoltan Trocsanyi, “The Dipole formalism for next-to-leading order QCD calculations with massive partons”, *Nucl. Phys. B*, vol. 627, pp. 189–265, 2002.
- [247] *Les Houches 2017: Physics at TeV Colliders Standard Model Working Group Report*, 3 2018.
- [248] George F. Sterman and Maria E. Tejeda-Yeomans, “Multiloop amplitudes and resummation”, *Phys. Lett. B*, vol. 552, pp. 48–56, 2003.
- [249] Thomas Becher and Matthias Neubert, “Infrared singularities of scattering amplitudes in perturbative QCD”, *Phys. Rev. Lett.*, vol. 102, pp. 162001, 2009, [Erratum: Phys.Rev.Lett. 111, 199905 (2013)].
- [250] Einan Gardi and Lorenzo Magnea, “Factorization constraints for soft anomalous dimensions in QCD scattering amplitudes”, *JHEP*, vol. 03, pp. 079, 2009.
- [251] William B. Kilgore and Christian Sturm, “Two-Loop Virtual Corrections to Drell-Yan Production at order $\alpha_s\alpha^3$ ”, *Phys. Rev. D*, vol. 85, pp. 033005, 2012.
- [252] William B. Kilgore, “The Two-Loop Infrared Structure of Amplitudes with Mixed Gauge Groups”, *Eur. Phys. J. C*, vol. 73, pp. 2603, 2013.
- [253] A. Mitov and S. Moch, “The Singular behavior of massive QCD amplitudes”, *JHEP*, vol. 05, pp. 001, 2007.
- [254] Thomas Becher and Kirill Melnikov, “Two-loop QED corrections to Bhabha scattering”, *JHEP*, vol. 06, pp. 084, 2007.
- [255] Thomas Becher and Matthias Neubert, “Infrared singularities of QCD amplitudes with massive partons”, *Phys. Rev. D*, vol. 79, pp. 125004, 2009, [Erratum: Phys.Rev.D 80, 109901 (2009)].
- [256] Taushif Ahmed, Johannes M. Henn, and Matthias Steinhauser, “High energy behaviour of form factors”, *JHEP*, vol. 06, pp. 125, 2017.
- [257] J. Blümlein, P. Marquard, and N. Rana, “Asymptotic behavior of the heavy quark form factors at higher order”, *Phys. Rev. D*, vol. 99, no. 1, pp. 016013, 2019.

- [258] Stefano Catani, Massimiliano Grazzini, and A. Torre, “Transverse-momentum resummation for heavy-quark hadroproduction”, *Nucl. Phys. B*, vol. 890, pp. 518–538, 2014.
- [259] Stefano Catani, Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, Javier Mazzitelli, and Hayk Sargsyan, “Top-quark pair hadroproduction at next-to-next-to-leading order in QCD”, *Phys. Rev. D*, vol. 99, no. 5, pp. 051501, 2019.
- [260] Stefano Catani, Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, and Javier Mazzitelli, “Top-quark pair production at the LHC: Fully differential QCD predictions at NNLO”, *JHEP*, vol. 07, pp. 100, 2019.
- [261] Stefano Catani, Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, and Javier Mazzitelli, “Bottom-quark production at hadron colliders: fully differential predictions in NNLO QCD”, *JHEP*, vol. 03, pp. 029, 2021.
- [262] Luca Buonocore, Massimiliano Grazzini, and Francesco Tramontano, “The q_T subtraction method: electroweak corrections and power suppressed contributions”, *Eur. Phys. J. C*, vol. 80, no. 3, pp. 254, 2020.
- [263] A. V. Kotikov, “Differential equations method: New technique for massive Feynman diagrams calculation”, *Phys. Lett. B*, vol. 254, pp. 158–164, 1991.
- [264] Ettore Remiddi, “Differential equations for Feynman graph amplitudes”, *Nuovo Cim. A*, vol. 110, pp. 1435–1452, 1997.
- [265] T. Gehrmann and E. Remiddi, “Differential equations for two loop four point functions”, *Nucl. Phys.*, vol. B580, pp. 485–518, 2000.
- [266] J. Ablinger, A. Behring, J. Blümlein, A. De Freitas, A. von Manteuffel, and C. Schneider, “Calculating Three Loop Ladder and V-Topologies for Massive Operator Matrix Elements by Computer Algebra”, *Comput. Phys. Commun.*, vol. 202, pp. 33–112, 2016.
- [267] J. Ablinger, J. Blümlein, P. Marquard, N. Rana, and C. Schneider, “Automated Solution of First Order Factorizable Systems of Differential Equations in One Variable”, *Nucl. Phys. B*, vol. 939, pp. 253–291, 2019.
- [268] A. Goncharov, “Polylogarithms in arithmetic and geometry”, *Proceedings of the International Congree of Mathematicians*, vol. 1,2, pp. 374–387, 1995.
- [269] A. B. Goncharov, “Multiple polylogarithms and mixed Tate motives”, 3 2001.
- [270] E. Remiddi and J.A.M. Vermaseren, “Harmonic polylogarithms”, *Int.J.Mod.Phys.*, vol. A15, pp. 725–754, 2000.
- [271] Jens Vollinga and Stefan Weinzierl, “Numerical evaluation of multiple polylogarithms”, *Comput. Phys. Commun.*, vol. 167, pp. 177, 2005.
- [272] P. A. Zyla et al., “Review of Particle Physics”, *PTEP*, vol. 2020, no. 8, pp. 083C01, 2020.
- [273] R. Bonciani, G. Degraffi, and A. Vicini, “On the Generalized Harmonic Polylogarithms of One Complex Variable”, *Comput. Phys. Commun.*, vol. 182, pp. 1253–1264, 2011.

- [274] Neelima Agarwal, Lorenzo Magnea, Chiara Signorile-Signorile, and Anurag Tripathi, “The infrared structure of perturbative gauge theories”, *Phys. Rept.*, vol. 994, pp. 1–120, 2023.
- [275] G.P. Korchemsky and A.V. Radyushkin, “Renormalization of the wilson loops beyond the leading order”, *Nuclear Physics B*, vol. 283, pp. 342–364, 1987.
- [276] R. Bonciani et al., “Two-Loop Four-Fermion Scattering Amplitude in QED”, *Phys. Rev. Lett.*, vol. 128, no. 2, pp. 022002, 2022.
- [277] Andrea Ferroglia, Matthias Neubert, Ben D. Pecjak, and Li Lin Yang, “Two-loop divergences of massive scattering amplitudes in non-abelian gauge theories”, *JHEP*, vol. 11, pp. 062, 2009.
- [278] A. Blondel, J. Gluza, S. Jadach, P. Janot, and T. Riemann, Eds., *Theory for the FCC-ee: Report on the 11th FCC-ee Workshop Theory and Experiments*, vol. 3/2020 of *CERN Yellow Reports: Monographs*, Geneva, 5 2019. CERN.
- [279] D. R. Yennie, Steven C. Frautschi, and H. Suura, “The infrared divergence phenomena and high-energy processes”, *Annals Phys.*, vol. 13, pp. 379–452, 1961.
- [280] Ayres Freitas and Qian Song, “Two-Loop Electroweak Corrections with Fermion Loops to $e+e-\rightarrow ZH$ ”, *Phys. Rev. Lett.*, vol. 130, no. 3, pp. 031801, 2023.
- [281] Xiang Chen, Xin Guan, Chuan-Qi He, Zhao Li, Xiao Liu, and Yan-Qing Ma, “Complete two-loop electroweak corrections to $e^+e^- \rightarrow HZ$ ”, 9 2022.
- [282] Iris Bree et al., “Geometric Bookkeeping Guide to Feynman Integral Reduction and ϵ -Factorized Differential Equations”, *Phys. Rev. Lett.*, vol. 136, no. 24, pp. 241602, 2026.
- [283] Fabio Cascioli, Philipp Maierhöfer, and Stefano Pozzorini, “Scattering Amplitudes with Open Loops”, *Phys. Rev. Lett.*, vol. 108, pp. 111601, 2012.
- [284] Federico Buccioni, Stefano Pozzorini, and Max Zoller, “On-the-fly reduction of open loops”, *Eur. Phys. J. C*, vol. 78, no. 1, pp. 70, 2018.
- [285] Ansgar Denner, Jean-Nicolas Lang, and Sandro Uccirati, “Recola2: REcursive Computation of One-Loop Amplitudes 2”, *Comput. Phys. Commun.*, vol. 224, pp. 346–361, 2018.
- [286] Ansgar Denner, Stefan Dittmaier, and Lars Hofer, “Collier: a fortran-based Complex One-Loop Library in Extended Regularizations”, *Comput. Phys. Commun.*, vol. 212, pp. 220–238, 2017.
- [287] S. Kallweit, J. M. Lindert, S. Pozzorini, and M. Schönherr, “NLO QCD+EW predictions for $2\ell 2\nu$ diboson signatures at the LHC”, *JHEP*, vol. 11, pp. 120, 2017.
- [288] Stefan Dittmaier, “A General approach to photon radiation off fermions”, *Nucl. Phys. B*, vol. 565, pp. 69–122, 2000.
- [289] Stefan Dittmaier, Alois Kabelschacht, and Tobias Kasprzik, “Polarized QED splittings of massive fermions and dipole subtraction for non-collinear-safe observables”, *Nucl. Phys. B*, vol. 800, pp. 146–189, 2008.
- [290] T. Gehrmann and N. Greiner, “Photon Radiation with MadDipole”, *JHEP*, vol. 12, pp. 050, 2010.

- [291] Marek Schönherr, “An automated subtraction of NLO EW infrared divergences”, *Eur. Phys. J. C*, vol. 78, no. 2, pp. 119, 2018.
- [292] Stefano Catani, Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, Javier Mazzitelli, and Chiara Savoini, “Higgs Boson Production in Association with a Top-Antitop Quark Pair in Next-to-Next-to-Leading Order QCD”, *Phys. Rev. Lett.*, vol. 130, no. 11, pp. 111902, 2023.
- [293] Luca Buonocore, Simone Devoto, Stefan Kallweit, Javier Mazzitelli, Luca Rottoli, and Chiara Savoini, “Associated production of a W boson and massive bottom quarks at next-to-next-to-leading order in QCD”, *Phys. Rev. D*, vol. 107, no. 7, pp. 074032, 2023.
- [294] Luca Buonocore, Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, Javier Mazzitelli, Luca Rottoli, and Chiara Savoini, “Precise Predictions for the Associated Production of a W Boson with a Top-Antitop Quark Pair at the LHC”, *Phys. Rev. Lett.*, vol. 131, no. 23, pp. 231901, 2023.
- [295] Simone Devoto, Massimiliano Grazzini, Stefan Kallweit, Javier Mazzitelli, and Chiara Savoini, “Precise predictions for $t\bar{t}H$ production at the LHC: inclusive cross section and differential distributions”, 11 2024.
- [296] Stefano Catani, Simone Devoto, Massimiliano Grazzini, and Javier Mazzitelli, “Soft-parton contributions to heavy-quark production at low transverse momentum”, *JHEP*, vol. 04, pp. 144, 2023.
- [297] Valerio Bertone, Stefano Carrazza, Nathan P. Hartland, and Juan Rojo, “Illuminating the photon content of the proton within a global PDF analysis”, *SciPost Phys.*, vol. 5, no. 1, pp. 008, 2018.
- [298] Gavin P. Salam and Emma Slade, “Cuts for two-body decays at colliders”, *JHEP*, vol. 11, pp. 220, 2021.
- [299] Richard D. Ball et al., “Determination of the theory uncertainties from missing higher orders on NNLO parton distributions with percent accuracy”, *Eur. Phys. J. C*, vol. 84, no. 5, pp. 517, 2024.
- [300] John C. Collins and Davison E. Soper, “Angular Distribution of Dileptons in High-Energy Hadron Collisions”, *Phys. Rev. D*, vol. 16, pp. 2219, 1977.
- [301] Albert M Sirunyan et al., “Search for resonant and nonresonant new phenomena in high-mass dilepton final states at $\sqrt{s} = 13$ TeV”, *JHEP*, vol. 07, pp. 208, 2021.
- [302] Elie Hammou, Zahari Kassabov, Maeve Madigan, Michelangelo L. Mangano, Luca Mantani, James Moore, Manuel Morales Alvarado, and Maria Ubiali, “Hide and seek: how PDFs can conceal new physics”, *JHEP*, vol. 11, pp. 090, 2023.
- [303] Simone Amoroso, Mauro Chiesa, Clara Lavinia Del Pio, Katerina Lipka, Fulvio Piccinini, Federico Vazzoler, and Alessandro Vicini, “Probing the weak mixing angle at high energies at the LHC and HL-LHC”, *Phys. Lett. B*, vol. 844, pp. 138103, 2023.
- [304] U. Baur, S. Keller, and W. K. Sakumoto, “QED radiative corrections to Z boson production and the forward backward asymmetry at hadron colliders”, *Phys. Rev. D*, vol. 57, pp. 199–215, 1998.
- [305] Albert M. Sirunyan et al., “Measurement of the weak mixing angle using the forward-backward asymmetry of Drell-Yan events in pp collisions at 8 TeV”, *Eur. Phys. J. C*, vol. 78, no. 9, pp. 701, 2018.

- [306] Georges Aad et al., “Measurement of the forward-backward asymmetry of electron and muon pair-production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”, *JHEP*, vol. 09, pp. 049, 2015.
- [307] Roel Aaij et al., “Measurement of the forward-backward asymmetry in $Z/\gamma^* \rightarrow \mu^+\mu^-$ decays and determination of the effective weak mixing angle”, *JHEP*, vol. 11, pp. 190, 2015.
- [308] “Measurement of the Drell-Yan forward-backward asymmetry and of the effective leptonic weak mixing angle using proton-proton collisions at 13 TeV”, 2024.
- [309] S. Navas et al., “Review of particle physics”, *Phys. Rev. D*, vol. 110, no. 3, pp. 030001, 2024.
- [310] Matteo Cacciari, Mario Greco, and Paolo Nason, “The P(T) spectrum in heavy flavor hadroproduction”, *JHEP*, vol. 05, pp. 007, 1998.
- [311] John C. Collins and Wu-Ki Tung, “Calculating Heavy Quark Distributions”, *Nucl. Phys. B*, vol. 278, pp. 934, 1986.
- [312] Tommaso Giani, Giacomo Magni, and Juan Rojo, “SMEFiT: a flexible toolbox for global interpretations of particle physics data with effective field theories”, *Eur. Phys. J. C*, vol. 83, no. 5, pp. 393, 2023.
- [313] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, “Importance Nested Sampling and the MultiNest Algorithm”, *Open J. Astrophys.*, vol. 2, no. 1, pp. 10, 2019.
- [314] Farhan Feroz and M. P. Hobson, “Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis”, *Mon. Not. Roy. Astron. Soc.*, vol. 384, pp. 449, 2008.
- [315] Riccardo Barbieri, Gino Isidori, Joel Jones-Perez, Paolo Lodone, and David M. Straub, “ $U(2)$ and Minimal Flavour Violation in Supersymmetry”, *Eur. Phys. J. C*, vol. 71, pp. 1725, 2011.
- [316] Riccardo Barbieri, Dario Buttazzo, Filippo Sala, and David M. Straub, “Flavour physics from an approximate $U(2)^3$ symmetry”, *JHEP*, vol. 07, pp. 181, 2012.
- [317] Jacob J. Ethier, Giacomo Magni, Fabio Maltoni, Luca Mantani, Emanuele R. Nocera, Juan Rojo, Emma Slade, Eleni Vryonidou, and Cen Zhang, “Combined SMEFT interpretation of Higgs, diboson, and top quark data from the LHC”, *JHEP*, vol. 11, pp. 089, 2021.
- [318] Ilaria Brivio, “SMEFTsim 3.0 — a practical guide”, *JHEP*, vol. 04, pp. 073, 2021.
- [319] Jaco ter Hoeve, Luca Mantani, Juan Rojo, Alejo N. Rossia, and Eleni Vryonidou, “Higgs trilinear coupling in the standard model effective field theory at the high luminosity LHC and the FCC-ee”, *Phys. Rev. D*, vol. 112, no. 1, pp. 013008, 2025.
- [320] Jaco ter Hoeve, Luca Mantani, Juan Rojo, Alejo N. Rossia, and Eleni Vryonidou, “Connecting scales: RGE effects in the SMEFT at the LHC and future colliders”, *JHEP*, vol. 06, pp. 125, 2025.
- [321] Georges Aad et al., “Differential cross-section measurements for the electroweak production of dijets in association with a Z boson in proton–proton collisions at ATLAS”, *Eur. Phys. J. C*, vol. 81, no. 2, pp. 163, 2021.

- [322] Admir Greljo, Shayan Iranipour, Zahari Kassabov, Maeve Madigan, James Moore, Juan Rojo, Maria Ubiali, and Cameron Voisey, “Parton distributions in the SMEFT from high-energy Drell-Yan tails”, *JHEP*, vol. 07, pp. 122, 2021.
- [323] Admir Greljo, Jakub Salko, Aleks Smolkovič, and Peter Stangl, “Rare b decays meet high-mass Drell-Yan”, *JHEP*, vol. 05, pp. 087, 2023.
- [324] Eamonn Maguire, Lukas Heinrich, and Graeme Watt, “HEPData: a repository for high energy physics data”, *J. Phys. Conf. Ser.*, vol. 898, no. 10, pp. 102006, 2017.
- [325] Céline Degrande, Gauthier Durieux, Fabio Maltoni, Ken Mimasu, Eleni Vryonidou, and Cen Zhang, “Automated one-loop computations in the standard model effective field theory”, *Phys. Rev. D*, vol. 103, no. 9, pp. 096024, 2021.
- [326] John Ellis, Maeve Madigan, Ken Mimasu, Veronica Sanz, and Tevong You, “Top, Higgs, Diboson and Electroweak Fit to the Standard Model Effective Field Theory”, *JHEP*, vol. 04, pp. 279, 2021.
- [327] et al. T. Virdee, “Lep3 projections”, 2025.
- [328] M. Benedikt et al., “Future Circular Collider Feasibility Study Report: Volume 1, Physics, Experiments, Detectors”, *Eur. Phys. J. C*, vol. 85, no. 12, pp. 1468, 2025.
- [329] M. Benedikt et al., “Future Circular Collider Feasibility Study Report: Volume 2, Accelerators, Technical Infrastructure and Safety”, *Eur. Phys. J. ST*, vol. 234, no. 19, pp. 5713–6197, 2025.
- [330] Eugenia Celada, Tommaso Giani, Jaco ter Hoeve, Luca Mantani, Juan Rojo, Alejo N. Rossia, Marion O. A. Thomas, and Eleni Vryonidou, “Mapping the SMEFT at high-energy colliders: from LEP and the (HL-)LHC to the FCC-ee”, *JHEP*, vol. 09, pp. 091, 2024.
- [331] “Prospects in electroweak, Higgs and Top physics at FCC”, 3 2025.
- [332] Alain BLONDEL, “Statistical and experimental systematic uncertainties on ewpos at the z pole in fcc-ee”, March 2025.
- [333] H. Abramowicz et al., “A Linear Collider Vision for the Future of Particle Physics”, 3 2025.
- [334] Alexander Aryshev et al., “The International Linear Collider: Report to Snowmass 2021”, 3 2022.
- [335] Jorge de Blas et al., “Physics Briefing Book: Input for the 2026 update of the European Strategy for Particle Physics”, 11 2025.
- [336] Hannes Mildner, “An EWPD SMEFT likelihood for the LHC — and how to improve it with measurements of W and Z boson properties”, *JHEP*, vol. 07, pp. 089, 2025.
- [337] A. Freitas et al., “Theoretical uncertainties for electroweak and Higgs-boson precision measurements at FCC-ee”, 6 2019.