



UNIVERSITÀ
DI PAVIA

DIPARTIMENTO DI FISICA
CORSO DI LAUREA IN SCIENZE FISICHE

Search for Charged Lepton Flavour Violation in $Z \rightarrow \ell\ell'$ decays at the CMS Experiment

Ricerca di violazione del sapore per leptoni carichi nei
decadimenti $Z \rightarrow \ell\ell'$ all'esperimento CMS

Candidata:
Valeria Farina

Supervisori:

Prof. Mario Pelliccioni

Prof.ssa Cristina Riccardi

Co-Supervisore:

Dott.ssa Chiara Aimè

A.A. 2025-2026

Abstract

Charged lepton flavour violation is strongly suppressed in the Standard Model, and its observation would provide unambiguous evidence of a new physics phenomenon. This thesis presents a search for the lepton flavour-violating decay $Z \rightarrow e\mu$ in proton–proton collisions at $\sqrt{s} = 13.6$ TeV, recorded by the CMS experiment at the LHC in 2024 and corresponding to an integrated luminosity of 109.95 fb^{-1} . A signal would appear as a narrow peak at the Z boson mass in the invariant mass distribution of opposite-charge electron–muon pairs. No significant excess over the expected background is observed. The observed (expected) upper limit on the branching fraction at 95% confidence level is $1.4 (2.0) \times 10^{-7}$, including statistical uncertainties only. The expected sensitivity is comparable to that of the previous CMS search, based on 138 fb^{-1} at 13 TeV, and the analysis provides the basis for its extension to the full dataset at 13.6 TeV.

Sommario

La violazione del sapore per leptoni carichi è fortemente soppressa nel Modello Standard e la sua osservazione costituirebbe un'evidenza inequivocabile di nuova fisica. Questa tesi presenta la ricerca del decadimento $Z \rightarrow e\mu$, che viola il sapore leptonico, nelle collisioni protone-protone a $\sqrt{s} = 13.6$ TeV registrate dall'esperimento CMS di LHC nel 2024, corrispondenti a una luminosità integrata di 109.95 fb^{-1} . Un segnale apparirebbe come un picco in corrispondenza della massa del bosone Z nella distribuzione di massa invariante delle coppie elettrone-muone di carica opposta. Non si osserva alcun eccesso significativo rispetto al fondo atteso. Il limite superiore osservato (atteso) sul branching fraction al 95% di livello di confidenza è $1.4 (2.0) \times 10^{-7}$, considerando le sole incertezze statistiche. La sensibilità attesa è paragonabile a quella della precedente ricerca di CMS, basata su 138 fb^{-1} a 13 TeV, e l'analisi costituisce la base per la ricerca nell'intero dataset raccolto a 13.6 TeV.

Contents

Abstract	iii
Sommario	v
Introduction	1
1 Theoretical overview	3
1.1 The Standard Model of particle physics	3
1.2 Charged Lepton Flavour Violation	6
1.3 Beyond the Standard Model	9
1.4 Experimental searches for CLFV	10
1.5 Z boson production	11
1.6 Analysis strategy	12
2 The Compact Muon Solenoid experiment	13
2.1 The Large Hadron Collider	13
2.1.1 LHC operating periods	16
2.1.2 LHC experiments	16
2.2 The Compact Muon Solenoid detector	17
2.2.1 Inner tracker	20
2.2.2 Calorimeters	21
2.2.3 Muon system	22
2.2.4 Trigger system	24
2.3 Physics object reconstruction and identification	25
2.3.1 Event reconstruction and particle flow algorithm	25
2.3.2 Electron reconstruction and identification	26
2.3.3 Complex objects reconstruction	28
2.3.4 Muon reconstruction and identification	30
2.4 CMS Run 3 data flow	32

3	Event selection	33
3.1	Data and Monte Carlo samples	35
3.1.1	Signal sample	36
3.1.2	Background samples	36
3.2	Selection algorithm	39
3.2.1	Trigger selection	39
3.2.2	Lepton selection	40
3.2.3	Event-level selection	44
3.2.4	Pair selection	47
3.2.5	Selection efficiency	51
3.2.6	MC and Data comparison	54
4	Signal and background modelling	57
4.1	Fitting strategy	57
4.2	Signal model	59
4.3	Background model	60
4.4	Total fit and validation	62
4.4.1	Combined model and branching fraction	63
4.4.2	Validation with pseudo-experiments	65
4.5	Fit to data	66
5	Statistical interpretation and results	71
5.1	Background-only hypothesis test	72
5.2	Upper limit	72
5.3	Comparison with the Run 2 result	75
	Conclusions	77
	Bibliography	79

Introduction

The Standard Model of particle physics provides a remarkably successful description of elementary particles and their interactions, yet it does not account for all the phenomena observed in nature. Neutrino oscillations demonstrate that lepton flavour is not conserved in the neutral sector, motivating searches for charged lepton flavour violation (CLFV). Although no CLFV process has been observed to date, searches for these rare decays provide a powerful probe of physics beyond the Standard Model.

At the Large Hadron Collider (LHC), the copious production of Z bosons makes their decays into two charged leptons of different flavours a natural place to look for CLFV. Among the possible final states, the electron–muon pair offers the cleanest signature: both leptons are directly reconstructed, and a signal would appear as a narrow peak in the invariant mass spectrum at the Z boson mass. The most stringent limit on this channel, 1.9×10^{-7} on the branching fraction at 95% confidence level (CL), was set by the Compact Muon Solenoid (CMS) Collaboration with 138 fb^{-1} of data at 13 TeV [1].

This thesis presents a search for $Z \rightarrow e\mu$ decay with the proton–proton collisions recorded by CMS in 2024 at 13.6 TeV, corresponding to an integrated luminosity of 109.95 fb^{-1} . Data are used to search for a signal and, in its absence, to set an upper limit on the branching fraction. The analysis also aims to establish a strategy applicable to the full dataset collected at 13.6 TeV, corresponding to an integrated luminosity of 350 fb^{-1} .

The search relies on a dedicated event selection and on a fit to the invariant mass distribution of the electron–muon pairs. No significant excess over the background expectation is found, and an upper limit of 1.4×10^{-7} (95% CL) is set on the branching fraction, considering statistical uncertainties only.

The thesis is organised as follows. Chapter 1 introduces the theoretical context, explaining why CLFV is suppressed in the Standard Model and why Z boson decays offer a sensitive probe of new physics phenomena. Chapter 2 describes the LHC and the CMS detector, focusing on the reconstruction and identification of

electrons and muons on which the search relies. Chapter 3 presents the collision data and simulated samples and the event selection, designed to isolate the $e\mu$ signature from the processes that can mimic it. Chapter 4 builds the model of the $e\mu$ invariant mass distribution, validates it on simulation and applies it to the unblinded data. Chapter 5 presents the statistical interpretation, through a test of the background-only hypothesis and an upper limit on the branching fraction. Finally, the Conclusions summarise the main results and discuss the prospects for extending the analysis to a larger dataset.

Chapter 1

Theoretical overview

The Standard Model of particle physics provides a remarkably successful description of elementary particles and their interactions. Only its main features relevant to this work are briefly reviewed in this chapter. More detailed discussions can be found in Refs. [2, 3].

Despite its success, the Standard Model is incomplete: an important exception is the neutrino sector. Neutrino oscillations imply neutral lepton flavour violation, motivating searches for charged lepton flavour violation, which has not yet been observed.

In the Standard Model extended to include neutrino masses, these processes are allowed but strongly suppressed, far below current experimental sensitivities. The observation of one of these processes would therefore constitute evidence of physics beyond the Standard Model.

1.1 The Standard Model of particle physics

The Standard Model (SM) of particle physics is a quantum field theory that provides a consistent description of the elementary particles and their interactions, combining the principles of quantum mechanics and special relativity. Within the SM, interactions are described by gauge symmetries and are mediated by gauge bosons, which are the force-carrying particles.

The SM is a gauge theory based on the unitary symmetry group

$$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y. \quad (1.1)$$

The three factors describe the strong and electroweak interactions. The first factor, $SU(3)_C$, is the gauge symmetry of Quantum Chromodynamics (QCD), the

theory of the strong interaction. The remaining two factors, $SU(2)_L \otimes U(1)_Y$, constitute the electroweak sector, which unifies the weak and electromagnetic interactions.

The number of gauge bosons associated with a symmetry group is determined by the number of its generators. For a special unitary group $SU(N)$, the number of generators is $N^2 - 1$, and the corresponding gauge theory contains $N^2 - 1$ gauge fields. The $SU(3)_C$ symmetry consequently gives rise to 8 gauge bosons, identified with the eight gluons mediating the strong interaction. Similarly, $SU(2)_L$ has 3 gauge bosons. The remaining factor, $U(1)_Y$, is an abelian group, it has a single generator and therefore one gauge boson. Thus, the electroweak sector contains four gauge fields.

Following the spontaneous symmetry breaking of $SU(2)_L \otimes U(1)_Y$ through the Higgs mechanism, these fields combine to form the physical W^+ , W^- , Z and photon (γ) bosons. The photon mediates the electromagnetic interaction, while the $W^\pm$ and Z bosons mediate the weak interaction.

The strong interaction described by $SU(3)_C$ is responsible for the dynamics of quarks and gluons and for the formation of colour-neutral hadrons, such as mesons and baryons. The electroweak gauge symmetry $SU(2)_L \otimes U(1)_Y$, on the other hand, is spontaneously broken at the electroweak scale down to the single residual symmetry $U(1)_{\text{EM}}$, associated with the electromagnetic interaction. This breaking is responsible for giving mass to the weak mediators, the $W^\pm$ and Z bosons, while the photon, associated with the unbroken $U(1)_{\text{EM}}$, remains massless. Gravity, the fourth fundamental interaction, is not included in the SM.

The elementary matter constituents of the SM are quarks and leptons, spin-1/2 fermions arranged into three generations, or families, as shown in Fig. 1.1, together with their corresponding antiparticles. The first generation forms ordinary matter, while the second and third generations contain heavier particles that are mostly unstable and decay into lighter ones.

Quarks carry colour charge, electric charge and weak isospin, and therefore participate in strong, electromagnetic and weak interactions. Charged leptons carry electric charge and weak isospin and therefore interact through both electromagnetic and weak interactions, while neutrinos have no electric charge and interact only through the weak interaction.

Each lepton generation is associated with a distinct flavour: electron, muon and tau. In the SM, charged lepton flavour is conserved in all gauge interactions. For




FIGURE 1.1: The SM elementary particle content. For each particle, its mass and quantum numbers are shown [4].

example, the charged-current interaction mediated by the W boson couples the electron to the electron neutrino, the muon to the muon neutrino, and the tau to the tau neutrino. This feature of the SM will be discussed in more detail in the following section.

While the electroweak theory successfully describes the observed interactions, explicit mass terms for the electroweak gauge bosons would break gauge invariance. However, there was experimental evidence, even before the direct discovery of the $W^\pm$ and Z bosons, that the mediators of the weak interaction had to be massive. In particular, the short timescale and corresponding short range of weak processes, such as nuclear beta decay, indicated that the weak interaction is mediated by massive particles. This was consistent with the expectation that a short range interaction is associated with a massive mediator. The $W^\pm$ and Z bosons were subsequently directly observed by the UA1 and UA2 experiments at CERN in 1983 [5, 6, 7, 8], confirming this picture. Their masses are generated within the SM through the Higgs mechanism [9, 10, 11], which resolves the apparent tension between massive gauge bosons and gauge invariance. The photon, in contrast,

remains massless. The same mechanism also generates the masses of the charged fermions through the Yukawa couplings.

The experimental discovery of the Higgs boson in 2012 by the Compact Muon Solenoid (CMS) and A Toroidal LHC ApparatuS (ATLAS) experiments of the Large Hadron Collider (LHC) at CERN [12, 13] provided the final confirmation of the SM particle content.

1.2 Charged Lepton Flavour Violation

In the SM, each lepton family is associated with a lepton flavour quantum number, namely L_e , L_μ , and L_τ , corresponding to the electron, muon, and tau generations, respectively. Charged leptons and their associated neutrinos carry one unit of the corresponding flavour number, while all other particles have zero lepton flavour. The assignments are summarized in Table 1.1.

Particle	L_e	L_μ	L_τ	$L = L_e + L_\mu + L_\tau$
e^-, ν_e	+1	0	0	+1
$e^+, \bar{\nu}_e$	-1	0	0	-1
μ^-, ν_μ	0	+1	0	+1
$\mu^+, \bar{\nu}_\mu$	0	-1	0	-1
τ^-, ν_τ	0	0	+1	+1
$\tau^+, \bar{\nu}_\tau$	0	0	-1	-1

TABLE 1.1: Lepton flavour numbers L_e , L_μ , L_τ and total lepton number L assigned to leptons and antileptons. All quarks and gauge bosons carry $L_e = L_\mu = L_\tau = 0$.

In the original formulation of the SM, where neutrinos are assumed to be massless, both total lepton number L and the individual lepton flavour numbers L_i , with $i = e, \mu, \tau$, are conserved quantities. This conservation emerges naturally from the SM and forbids transitions between leptons of different flavours.


This differs from the quark sector, where charged-current weak interactions mix quark flavours through the Cabibbo–Kobayashi–Maskawa (CKM) matrix, allowing flavour-changing transitions between different generations.

The observation of neutrino oscillations by experiments such as Super-Kamiokande [14] and SNO (Sudbury Neutrino Observatory) [15], demonstrates that neutrinos have

mass and that lepton flavour is not conserved in the neutral sector. This phenomenon is described by the mixing of neutrino flavour and mass eigenstates through the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix.

However, no analogous mixing has been observed among charged leptons. Charged Lepton Flavour Violation (CLFV) has never been detected, making it one of the most sensitive probes of physics beyond the SM.

In a minimal extension of the SM including non-zero neutrino masses, CLFV decays of the Z boson, generically denoted as $Z \rightarrow \ell_i^\pm \ell_j^\mp$ ($i \neq j$), are induced only through one-loop diagrams involving neutrino oscillations and W bosons, as illustrated in Fig. 1.2a.



(a) SM contribution: W -boson loop with neutrino oscillation.

(b) Effective vertex.

FIGURE 1.2: Feynman diagrams for the process $Z \rightarrow \ell_i \ell_j$. The effective vertex (b) encodes both the SM contribution, mediated by a W boson loop with neutrino oscillations shown explicitly in (a), and possible model independent beyond SM contributions.

Since the weak-interaction eigenstates, ν_e, ν_μ, ν_τ , are related to the neutrino mass eigenstates, ν_1, ν_2, ν_3 , through the PMNS matrix, the total amplitude is obtained by summing the contributions of the three neutrino mass eigenstates propagating in the loop. Because the PMNS matrix is unitary, all flavour-independent contributions cancel when the neutrino masses are equal. This cancellation, known as the leptonic analogue of the Glashow–Iliopoulos–Maiani (GIM) mechanism, implies that only the small differences between the neutrino masses survive in the amplitude. As a result,

$$\mathcal{A}(Z \rightarrow \ell_i \ell_j^1) \propto \frac{\Delta m_{\nu_{ij}}^2}{M_W^2} \quad i \neq j. \quad (1.2)$$

¹The lepton charges are omitted when not necessary.

The probability of a particle decaying through a specific final state is quantified by its branching fraction, defined as the ratio between the corresponding partial decay width and the total decay width,

$$\mathcal{B}(X \rightarrow f) = \frac{\Gamma(X \rightarrow f)}{\Gamma_{\text{tot}}}, \quad (1.3)$$

where $\Gamma(X \rightarrow f)$ is the partial decay width for the decay into the final state f , while Γ_{tot} is the total decay width of the particle X . The branching fraction therefore represents the fraction of decays occurring through a given final state. Since the partial width is proportional to the squared amplitude, $\Gamma(Z \rightarrow \ell_i \ell_j) \propto |\mathcal{A}|^2$, and the total width of the Z boson is dominated by decay modes already allowed in the SM, Γ_{tot} can be regarded as unaffected by the CLFV channel. The branching fraction thus scales with the squared amplitude, leading to

$$\mathcal{B}(Z \rightarrow \ell_i \ell_j) \propto \left(\frac{\Delta m_{\nu_{ij}}^2}{M_W^2} \right)^2 \quad i \neq j. \quad (1.4)$$

Since the measured neutrino mass-squared differences are many orders of magnitude smaller than the electroweak scale, CLFV branching fractions are predicted to be extraordinarily suppressed, far beyond the current experimental sensitivity. The resulting branching ratio predictions for the three flavour-violating Z boson decay channels in the minimally extended SM are summarized in Table 1.2 [16].

Channel	Predicted $\mathcal{B}$
$Z \rightarrow \mu^\mp \tau^\pm$	$\sim 10^{-54}$
$Z \rightarrow e^\mp \tau^\pm$	$\lesssim 4 \times 10^{-60}$
$Z \rightarrow e^\mp \mu^\pm$	$\lesssim 4 \times 10^{-60}$

TABLE 1.2: Predicted branching ratios for the three lepton flavour violating Z decay channels in the SM minimally extended with massive neutrinos [16].

Any observation of CLFV in Z decays at a level accessible to current experiments (branching fractions down to $\mathcal{O}(10^{-7})$ – $\mathcal{O}(10^{-9})$) would therefore be unambiguous evidence of physics beyond this minimal extension of the SM.

1.3 Beyond the Standard Model

Although the SM has been remarkably successful in describing elementary particles and their interactions, it is widely regarded as an effective theory valid up to some higher energy scale. Several experimental observations indicate that it is incomplete, including the existence of neutrino masses and oscillations, the necessity of dark matter, the matter–antimatter asymmetry of the Universe, and the absence of a consistent quantum theory of gravity. These open questions strongly motivate the search for physics Beyond the Standard Model (BSM).

CLFV provides one of the cleanest probes of BSM physics. Since CLFV is suppressed to unobservable levels in the minimal extension of the SM, the observation of any CLFV process would constitute unambiguous evidence for new physics. In many BSM scenarios, new particles or interactions naturally induce flavour changing couplings involving charged leptons, enhancing the branching fractions of CLFV decays by many orders of magnitude with respect to the SM prediction. Rather than considering a specific theoretical framework, these effects can be described in a model-independent way through an effective $Z \rightarrow \ell_i \ell_j$ vertex, represented in Fig. 1.2b. The effective interaction parametrizes the contribution of heavy virtual particles without specifying their underlying dynamics, providing a generic description of CLFV phenomena.

Such effective couplings arise in a broad class of BSM models, including supersymmetric theories, seesaw scenarios with heavy Majorana neutrinos, models with extended Higgs sectors, leptoquarks, and theories with the additional neutral gauge bosons Z' [1]. Depending on the masses and couplings of the new particles, these models can predict branching fractions for $Z \rightarrow \ell_i \ell_j$ decays many orders of magnitude larger than the SM expectation, in some cases within the reach of current or future collider experiments.

Because these predictions are highly model dependent, experimental searches can be interpreted in a model independent framework, with upper limits on the branching fractions that can subsequently be translated into constraints on a wide range of BSM scenarios.

1.4 Experimental searches for CLFV

Historically, the search for CLFV began with the $\mu \rightarrow e\gamma$ decay, first investigated by Hincks and Pontecorvo in 1947. Since then, increasingly sensitive searches for processes such as $\mu \rightarrow e\gamma$, $\mu \rightarrow 3e$, $\mu N \rightarrow eN$, and τ decays such as $\tau \rightarrow \mu\gamma$ have established some of the strongest experimental constraints on CLFV interactions. Owing to their exceptional precision, these low energy experiments probe energy scales well beyond the reach of present collider experiments. Indirect limits from low energy measurements can be translated into stringent constraints on flavour-violating Z boson decays, reaching $\mathcal{B}(Z \rightarrow e\mu) \leq 10^{-12}$, $\mathcal{B}(Z \rightarrow e\tau) \leq 10^{-7}$, and $\mathcal{B}(Z \rightarrow \mu\tau) \leq 2 \times 10^{-7}$ [17]. However, these bounds are inherently model dependent, as they rely on assumptions about the underlying CLFV interaction, such as the presence of very heavy mediating particles. Direct searches for $Z \rightarrow \ell_i \ell_j$ decays, on the other hand, probe this flavour violating vertex without relying on such assumptions. They therefore provide complementary information to low energy experiments, including ongoing and future facilities such as MEG II, Mu3e, DeeMe, Mu2e, and COMET [1].

Following the discovery of the Z boson, dedicated searches for CLFV Z boson decays have been carried out, evolving from the clean environment of leptonic collisions at the Large Electron Positron collider (LEP) to the large datasets collected by the LHC experiments.

LEP provided one of the first experimental constraints on the branching fractions $\mathcal{B}(Z \rightarrow e\mu)$, $\mathcal{B}(Z \rightarrow e\tau)$, and $\mathcal{B}(Z \rightarrow \mu\tau)$, reaching sensitivities at the $\mathcal{O}(10^{-5} - 10^{-6})$ level, as summarized in Tab. 1.3.

The larger production cross section of Z bosons at the LHC enabled a significant improvement in sensitivity. Using the full Run 2 dataset, corresponding to an integrated luminosity of 139 fb^{-1} , the ATLAS Collaboration performed searches for the $Z \rightarrow e\mu$, $Z \rightarrow e\tau$, and $Z \rightarrow \mu\tau$ decay modes, setting the most stringent limits on the channels involving τ leptons [18]. The CMS Collaboration searched for the $Z \rightarrow e\mu$ decay using the full Run 2 dataset, corresponding to an integrated luminosity of 138 fb^{-1} . The $e\mu$ final state provides a particularly clean experimental signature, allowing a highly sensitive search for CLFV. No significant excess above the SM expectation was observed, resulting in the current best direct limit on the branching fraction of this decay [1]. The current best experimental limits are reported in Table 1.3.

Experiment	Channel	Dataset	Upper limit (95% CL)
LEP	$Z \rightarrow e\mu$	Full LEP	$\sim 10^{-6}$
LEP	$Z \rightarrow e\tau$	Full LEP	$\sim 10^{-5}$
LEP	$Z \rightarrow \mu\tau$	Full LEP	$\sim 10^{-5}$
CMS	$Z \rightarrow e\mu$	Run 2	1.9×10^{-7}
ATLAS	$Z \rightarrow e\tau$	Run 2	5.0×10^{-6}
ATLAS	$Z \rightarrow \mu\tau$	Run 2	6.5×10^{-6}

TABLE 1.3: Evolution of the experimental sensitivity to CLFV Z boson decays. The current best limits are taken from [19].

1.5 Z boson production

At a proton–proton collider, Z bosons are produced through interactions between the partons, quarks and gluons, contained in the colliding protons. The dominant production mechanism is the Drell–Yan process,

$$q\bar{q} \rightarrow Z,$$

in which a quark and an antiquark annihilate to form a Z boson, as shown in the Feynman diagram in Fig. 1.3.




FIGURE 1.3: Leading-order Drell–Yan production of a Z boson.

At higher orders, additional partons can be emitted in the initial state, giving rise to processes such as

$$q\bar{q} \rightarrow Zg, \quad qg \rightarrow Zq,$$

which contribute to Z production accompanied by jets.

The large Z production rate at the Large Hadron Collider is particularly valuable for searches for rare decay modes. The inclusive Z production cross section at $\sqrt{s} = 13.6$ TeV is of the order of 6×10^4 pb [20], which, for an integrated luminosity

of $\mathcal{O}(100) \text{ fb}^{-1}$, yields to several billion Z bosons. This makes it possible to probe rare decays such as $Z \rightarrow \ell_i \ell_j$ with $i \neq j$.

1.6 Analysis strategy

This thesis presents a search for the decay of the Z boson in the CLFV channel $Z \rightarrow e\mu$. The main backgrounds are Drell–Yan induced processes, diboson and $t\bar{t}$ pair production.

The analysis is based on data collected by the CMS experiment during Run 3 (2022–2026), at a centre-of-mass energy of 13.6 TeV [20] and approximately 350 fb^{-1} of integrated luminosity [21]. In the present work, only the 2024 dataset is analyzed, corresponding to an integrated luminosity of 109.95 fb^{-1} .

In this work, I develop an analysis strategy in order to search for the CLFV decay $Z \rightarrow e\mu$. This analysis strategy aims at identifying signal candidates and suppressing SM background processes that can mimic the signal signature.

Since the signal consists of an electron and a muon originating from the decay of an on-shell Z boson, the analysis aims to identify events containing an opposite sign $e\mu$ pair with an invariant mass compatible with the Z boson mass. Candidate events are selected with criteria that enhance the signal sensitivity while suppressing SM background processes through suitable kinematic requirements.

The search is performed by studying the invariant mass distribution of the $e\mu$ system. A CLFV signal would appear as a narrow resonance around the Z boson mass. Since the SM expectation is negligible, any statistically significant excess would constitute evidence for CLFV and therefore for BSM physics.

A statistical analysis of the invariant mass spectrum is used to derive an upper limit on the branching fraction $\mathcal{B}(Z \rightarrow e\mu)$.

Chapter 2

The Compact Muon Solenoid experiment

The Compact Muon Solenoid is one of the four major experiments operating at the Large Hadron Collider. It played a central role in the discovery of the Higgs boson in 2012 and continues to explore both the SM and possible signatures of physics beyond it.

After introducing the main characteristics of the accelerator and its operating conditions, the detector is described with particular emphasis on the subsystems relevant for this analysis. The chapter also presents the reconstruction and identification of physics objects, including the particle-flow algorithm, and summarizes the data processing chain from data acquisition to the NanoAOD format used in this analysis.

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is the world's largest and highest energy particle accelerator. It was approved by the CERN (Conseil Européen pour la Recherche Nucléaire) Council at the end of 1994 and built from 1998 to 2008, aiming at the discovery of the Higgs boson and the study of rare processes. It consists of a circular underground tunnel with a circumference of approximately 27 km and is designed to operate at a nominal centre-of-mass energy of 14 TeV.

The choice of a proton–proton (pp) machine was motivated by several considerations. First of all, hadrons allow much higher centre-of-mass energy to be reached than electrons because the energy loss due to synchrotron radiation is much smaller. Furthermore, since protons are not elementary particles, the energy is shared among their constituent partons (gluons and quarks), allowing a wider

range of energies to be explored with fixed energy beams. Moreover, the large production cross sections associated with strong interactions make proton collisions particularly suitable for exploring rare processes and discovering new heavy particles. Finally, pp colliders can achieve higher luminosities than proton–antiproton colliders. For these reasons, the two key performance parameters of the LHC are the centre-of-mass energy and the luminosity. The nominal parameters of the LHC are summarised in Table 2.1.

Parameter	Nominal value
Beams	pp
Circumference	26.7 km
Centre-of-mass energy	14 TeV
Number of bunches	2808
Protons per bunch	10^{11}
Bunch crossing interval	25 ns
Instantaneous luminosity	$1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$

TABLE 2.1: Nominal parameters of the LHC [22].

Since protons are composite particles, only a fraction x of the beam energy is carried by the interacting partons. As a consequence, the partonic centre-of-mass energy is given by

$$\sqrt{s'} = \sqrt{x_1 x_2 s}$$

where x_1 and x_2 are the momentum fractions of the two interacting partons. Before being injected into the LHC, proton beams are progressively accelerated through the CERN accelerator complex (Fig. 2.1). Protons produced by the linear accelerator (LINAC4) are successively accelerated in the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS). In the PS, the proton beam is divided into bunches separated by 25 ns, which are then further accelerated in the SPS before being injected as two opposite beams into the LHC ring.

The instantaneous luminosity ($\mathcal{L}$) for two Gaussian beams colliding head-on, usually measured in $\text{cm}^{-2}\text{s}^{-1}$, can be calculated as:

$$\mathcal{L} = \frac{N_1 N_2}{4\pi\sigma_{xy}} f \quad (2.1)$$




FIGURE 2.1: The CERN accelerator complex [23]

with N_1 and N_2 the numbers of particles per bunch, σ_{xy} the transverse beam size, and f the bunch-crossing frequency.

The importance of the luminosity lies in its direct relation to the interaction rate, namely the number of events produced per unit time,

$$R = \frac{dN}{dt} = \mathcal{L}\sigma, \quad (2.2)$$

where σ is the cross section of the process. The cross section represents the probability for a given interaction to occur and has the dimensions of an area. It is commonly expressed in barn, where $1 \text{ b} = 10^{-24} \text{ cm}^2$.

Integrating the instantaneous luminosity over the data-taking period gives the integrated luminosity, usually expressed in fb^{-1} :

$$L_{\text{int}} = \int_0^T \mathcal{L}(t) dt. \quad (2.3)$$

The expected number of events produced for a process with cross section σ is therefore

$$N = L_{\text{int}}\sigma. \quad (2.4)$$

Each of the 2808 bunches circulating in the LHC ring contains approximately 10^{11} protons. Adjacent bunches cross every 25 ns at the interaction points. Since several

pp interactions may occur within the same bunch crossing, multiple interaction vertices are reconstructed in a single event. This phenomenon is known as pile-up [2, 22].

2.1.1 LHC operating periods

The operation of the LHC is organized into successive data-taking periods, known as Runs, which are separated by Long Shutdowns (LS). During these shutdowns, the accelerator complex and the experiments undergo maintenance, consolidation, and major upgrade activities to improve their performance and prepare for subsequent operation. Each Run is characterized by specific operating conditions, such as the centre-of-mass energy, the instantaneous luminosity, and the detector configuration.

Figure 2.2 summarizes the operating schedule of the LHC, from the first Run to the future High-Luminosity LHC (HL-LHC) programme. Run 1 (2010–2012) operated at centre-of-mass energies of 7 TeV and 8 TeV, while Run 2 (2015–2018) delivered pp collisions at 13 TeV. Following the completion of LS2, Run 3 started in 2022 at an increased centre-of-mass energy of 13.6 TeV and continued until 2026 [21]. After LS3, the HL-LHC programme will begin, aiming to increase the integrated luminosity by approximately one order of magnitude with respect to the LHC design value, thereby significantly extending the physics reach of the experiments [24].

The analysis presented in this thesis is based on pp collision data collected by CMS during Run 3, at a centre-of-mass energy of 13.6 TeV. Run 3 is expected to provide a total integrated luminosity of approximately 350 fb^{-1} by the end of the data-taking period in 2026, as reported in Fig. 2.3 [21].

2.1.2 LHC experiments

Four major experiments are located at the interaction points of the LHC. ATLAS (A Toroidal LHC ApparatuS) [26] and the Compact Muon Solenoid (CMS) [27] are general-purpose detectors, designed to investigate a broad range of physics scenarios, from precision measurements of the SM processes and Higgs boson properties to BSM physics. The CMS detector, which is the focus of this thesis, is described in detail in Section 2.2.




FIGURE 2.2: LHC project schedule [25]

The Large Hadron Collider beauty (LHCb) [28] experiment is dedicated to the study of hadrons containing beauty and charm quarks, with particular emphasis on flavour physics and the violation of the combined charge-conjugation and parity symmetry (CP violation). The A Large Ion Collider Experiment (ALICE) [29] is specifically designed to investigate strongly interacting matter at extremely high energy densities through heavy-ion collisions, providing insight into the properties of the quark-gluon plasma.

2.2 The Compact Muon Solenoid detector

The Compact Muon Solenoid (CMS) is one of the two general-purpose experiments operating at the LHC. It was designed to investigate a broad range of physics phenomena, ranging from precision measurements of SM processes to searches for BSM physics.

The detector has a cylindrical geometry with a central barrel and two endcap sections, providing an almost hermetic coverage around the interaction point. A schematic view of the CMS detector is shown in Fig. 2.4.

The central feature of CMS is a superconducting solenoid with an internal diameter of 6 m, providing a uniform magnetic field of 3.8 T. The magnetic field bends




FIGURE 2.3: Run 3 total integrated luminosity [21]

the trajectories of charged particles, allowing precise momentum measurements. The detector follows the typical layered (“onion-like”) structure of modern collider experiments. Moving radially outward from the interaction point, it is composed of:

1. the silicon tracking system;
2. the electromagnetic and hadronic calorimeters;
3. the muon detection system.

The tracker and calorimeters are located inside the superconducting solenoid, while the muon chambers are embedded in the steel return yoke surrounding the magnet. In addition, forward calorimeters extend the detector acceptance at high pseudo-rapidities. The main features of these subdetectors are described in the following sections [27].

Before discussing the individual detector components, it is useful to introduce the CMS coordinate system, shown in Fig. 2.5.

A right-handed Cartesian coordinate system is adopted, with its origin at the nominal interaction point. The x -axis points towards the centre of the LHC ring, the y -axis points vertically upward, and the z -axis is defined along the beam direction,




FIGURE 2.4: The CMS detector [30].

pointing towards the Jura mountains. Cylindrical coordinates are commonly used, where $r = \sqrt{x^2 + y^2}$ is the radial distance from the beam axis, ϕ is the azimuthal angle in the transverse plane, and θ is the polar angle measured with respect to the z -axis.

Two kinematic quantities are widely used throughout this thesis: the transverse momentum, defined as

$$p_T = p \sin \theta, \quad (2.5)$$

where p is the magnitude of the particle momentum, and the pseudorapidity, defined as

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right], \quad (2.6)$$

which describes the direction of a particle with respect to the beam axis. Particles produced perpendicular to the beam line have $\eta = 0$, whereas particles emitted close to the beam direction correspond to large values of $|\eta|$ [27].




FIGURE 2.5: The CMS coordinate system [31].

2.2.1 Inner tracker

The tracker is the innermost subdetector of the CMS experiment, surrounding the interaction point. Its primary purpose is the reconstruction of charged particle trajectories and interaction vertices with high precision. It is entirely based on silicon technology, has a length of approximately 5.8 m and a diameter of 2.5 m, and provides coverage up to pseudorapidities of about $|\eta| < 2.5$.

Immersed in the 3.8 T magnetic field generated by the superconducting solenoid, the tracker measures the curvature of charged-particle trajectories, allowing the determination of their transverse momentum according to

$$p_T = 0.3BR [\text{GeV}], \quad (2.7)$$

where B is the magnetic field strength, expressed in Tesla, R is the radius of curvature of the particle trajectory, measured in meters, and the transverse momentum is obtained in GeV.

The tracker consists of two complementary subsystems: the silicon pixel detector and the silicon strip detector. The pixel detector surrounds the interaction region and provides the first high-precision measurements of charged-particle trajectories. Its excellent spatial resolution enables the accurate reconstruction of the primary vertex, defined as the pp interaction point of the hard-scattering process, as well as of secondary vertices originating from the decays of long lived particles. It also provides precise measurements of the track impact parameters, defined as the distance of closest approach between a reconstructed track and a reference vertex (typically the primary vertex) in the transverse (d_{xy}) and longitudinal (d_z)

directions, which play an important role in the identification of displaced particles and in b -tagging algorithms.

The outer region is instrumented with silicon strip detectors, providing additional position measurements farther from the interaction point, which improves the determination of the track curvature and therefore the momentum resolution [32]. A major upgrade of the CMS tracker is currently planned for the HL-LHC era. The new all-silicon tracker will extend the pseudorapidity coverage and improve the detector performance under the significantly higher pile-up conditions expected after LS3 [33].

2.2.2 Calorimeters

Outside the tracker, the CMS calorimeter system is composed of two detectors: the electromagnetic calorimeter (ECAL) and the hadron calorimeter (HCAL). Their purpose is to measure the energy deposited by particles through the development of electromagnetic and hadronic showers. Together, they provide energy measurements for electrons, photons and hadrons, and play a key role in the reconstruction of jets and missing transverse energy (MET).

The ECAL is a homogeneous calorimeter made of lead tungstate (PbWO_4) crystals, chosen for their high density, fast response and excellent energy resolution. It is divided into a barrel section (EB), covering the region $|\eta| < 1.48$, and two endcaps (EE), extending the coverage up to approximately $|\eta| < 3.0$. A preshower detector is installed in front of the endcaps to improve the identification of photons and their separation from neutral pion (π_0) decays.

The PbWO_4 crystals are characterized by a short radiation length ($X_0 = 0.89$ cm), defined as the mean distance over which a high energy electron loses all but $1/e$ of its energy through bremsstrahlung, and by a small Molière radius ($R_M = 2.2$ cm), corresponding to the radius containing about 90% of the energy deposited by an electromagnetic shower. These properties result in compact showers, allowing a highly granular detector with excellent spatial and energy resolution [34].

The energy resolution of the ECAL can be parameterized as:

$$\left(\frac{\sigma(E)}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2,$$

where S is the stochastic term, accounting for statistical fluctuations in the shower development and scintillation light collection, N is the noise term, which includes

contributions from the detector electronics and pile-up, and C is the constant term, mainly arising from calibration uncertainties, detector non-uniformities and long-term stability. In the ECAL barrel, the stochastic and constant terms have been measured to be $S \approx 2.8\%$ and $C \approx 0.3\%$, using electrons in the energy range 20 GeV to 250 GeV. The noise term N is typically negligible at the high electron energies relevant for this analysis. The excellent energy resolution of the ECAL is essential for the precise reconstruction and identification of electrons [35].

The HCAL surrounds the ECAL and measures the energy of hadrons that are only partially absorbed in the electromagnetic calorimeter. It is composed of alternating layers of dense absorber material and plastic scintillators and is divided into a barrel (HB), endcap (HE), outer (HO, acting as a tail catcher for late-developing showers beyond the HB) and forward (HF) sections, providing nearly hermetic coverage up to $|\eta| < 5.2$. Combined with the ECAL, the HCAL contributes to the reconstruction of jets and to the measurement of the MET in the event [36].

For the HL-LHC era, the endcap calorimeters will be replaced by the High Granularity Calorimeter (HGCAL), a high precision sampling calorimeter designed to cope with the significantly increased radiation levels and pile-up conditions expected after LS3 [37].

2.2.3 Muon system

Muons are minimum-ionizing particles and lose only a small fraction of their energy while traversing the tracker and the calorimeters. As a consequence, they reach the outermost subdetector of CMS, the muon system, which is specifically designed for muon identification, triggering and momentum measurement. The momentum measurement is obtained by combining the information from the inner tracker with that from the muon chambers, providing an improved resolution over a wide momentum range.

A longitudinal view of one quarter of the CMS muon system is shown in Fig. 2.6. As illustrated in the figure, the muon system is embedded in the steel return yoke of the superconducting solenoid and consists of several layers of gaseous detectors interleaved with the absorber material. Different detector technologies are employed according to the expected particle rates and radiation levels in the various detector regions. Drift Tubes (DT) are used in the barrel region ($|\eta| < 1.2$), where the particle flux is relatively low, while Cathode Strip Chambers (CSC)




FIGURE 2.6: Longitudinal R-z view of one quadrant of the CMS muon system, including the Phase 2 upgrades planned for the HL-LHC. The different muon detector technologies are shown in the barrel and endcap regions. The interaction point is located at the lower left corner, while the dark grey regions represent the magnet return yoke [38].

are installed in the endcaps ($0.9 < |\eta| < 2.4$), where they provide excellent spatial resolution under higher background conditions. Resistive Plate Chambers (RPC) complement both systems by providing excellent time resolution and a fast trigger response. In addition, Gas Electron Multiplier (GEM) detectors are used in the forward region to provide additional measurements for muon reconstruction and triggering [38, 39].

The momentum resolution is affected by multiple scattering at low transverse momentum, which causes deviations of the particle trajectory as it traverses the detector material. At higher transverse momentum, multiple scattering becomes less significant, while the measurement of the track curvature becomes increasingly challenging. In this regime, the additional measurements provided by the muon chambers become particularly important for the momentum reconstruction.

The figure also illustrates the additional detector technologies foreseen for the upgrade of the CMS muon system for the HL-LHC. These include additional GEM detectors and improved RPC (iRPC), which will extend the muon detection capabilities towards the forward region and provide additional redundancy for muon triggering and reconstruction under the higher radiation and pile-up conditions

expected at the HL-LHC [40].

The different detector technologies provide complementary measurements, ensuring efficient muon identification and triggering over a wide pseudorapidity range.

2.2.4 Trigger system

At the LHC, proton bunches cross each other every 25 ns, corresponding to a bunch crossing frequency of 40 MHz. Since the size of a single collision event is of the order of ~ 1 MB, it is not possible to read out and store all the events produced by the detector. A trigger system is therefore required to select potentially interesting events and reduce the event rate to a level that can be stored and processed offline. The CMS trigger system is organized in two stages: the Level 1 Trigger (L1) and the High Level Trigger (HLT) [41].

The L1 trigger is a hardware based system that reduces the event rate from 40 MHz to approximately 100 kHz, with a latency of about 4 μ s. The stringent latency requirements limit the information available to L1, which relies primarily on the calorimeter and muon systems. At this stage, the full track reconstruction performed by the inner tracker cannot be used. Energy deposits in the ECAL and HCAL and signals from the muon detectors are processed into Trigger Primitives (TP), which are then combined to identify candidate objects and determine whether an event satisfies the requirements for further processing. The information from the calorimeter and muon trigger systems is finally combined by the Global Trigger, which makes the L1 decision [42].

Events accepted by L1 are subsequently processed by the HLT, a software-based system running on a computer farm. The HLT has access to the full detector information and performs a more detailed event reconstruction, allowing the event rate to be reduced further from approximately 100 kHz to approximately 1 kHz, which corresponds to the final data acquisition rate for offline analysis [43].

2.3 Physics object reconstruction and identification

The signals recorded by the CMS detector, consisting of hits in the tracker and muon chambers and energy deposits in the calorimeters, are processed to reconstruct and identify the particles produced in the collision. The particle flow algorithm combines the information from the different detector subsystems to reconstruct individual particles in the event. The reconstruction and identification of the physics objects relevant for the $Z \rightarrow e\mu$ analysis are described in the following sections.

2.3.1 Event reconstruction and particle flow algorithm

The reconstruction of physics objects starts from the information collected by the different CMS subdetectors. The Particle Flow (PF) algorithm combines the measurements from the tracker, calorimeters and muon system to reconstruct and identify the individual particles produced in the collision. The reconstruction proceeds from basic detector elements, such as charged particle tracks and calorimeter clusters, which are subsequently linked to form PF candidates.

The reconstruction of charged particle tracks is based on an iterative procedure designed to achieve high efficiency while keeping the rate of misreconstructed tracks low. The procedure starts from seeds, formed from combinations of hits in the tracker, which provide an initial estimate of the particle trajectory. Track candidates are then propagated through the detector using a Kalman filter, a recursive algorithm that combines the information from the hits already associated with the track with the measurement from the next detector layer to update the estimate of the particle trajectory and its associated uncertainties. If a compatible hit is found, it is incorporated into the track candidate; otherwise, the corresponding detector layer can be skipped. The procedure is repeated as the candidate is propagated through the tracker, and the associated hits are finally used to obtain the best estimate of the track parameters. Reconstructed tracks are then selected according to quality criteria based on the track fit and on their compatibility with the interaction region.

The reconstructed tracks are also used to determine the position of the collision vertices. Since multiple pp interactions can occur within the same bunch crossing,

several vertices can be reconstructed in a single event. The primary vertex is identified among them using the properties of the associated tracks, while the remaining vertices correspond to pile-up interactions [32].

In parallel, calorimeter clusters are reconstructed from energy deposits in the different calorimeter subsystems. The clustering procedure groups neighbouring cells around energy seeds, allowing the energy and direction of electromagnetic and hadronic activity to be measured and providing the basis for the reconstruction of neutral particles and electrons. In the ECAL, the clustering also accounts for the energy radiated by electrons through bremsstrahlung in the tracker material.

The reconstructed tracks and calorimeter clusters, together with information from the muon system, are then linked by the PF algorithm. The resulting PF candidates are classified as photons, electrons, muons, charged hadrons and neutral hadrons.

Electrons are reconstructed by associating a charged-particle track, obtained through a dedicated tracking algorithm that accounts for bremsstrahlung energy loss, with the corresponding ECAL energy deposits. Muons are reconstructed by combining information from the inner tracker and the muon system.

The PF candidates provide a particle-level description of the event and form the basis for the reconstruction of higher level physics objects, such as jets and missing transverse momentum [44]. The identification and isolation criteria applied to electrons and muons in the analysis are described in the following sections.

2.3.2 Electron reconstruction and identification

The reconstruction of electrons combines the information from the inner tracker and the ECAL. The electron trajectory is reconstructed using a Gaussian Sum Filter (GSF), which accounts for the non-Gaussian energy loss due to bremsstrahlung as the electron traverses the tracker material. The emitted photons deposit their energy in the ECAL, where the corresponding deposits are clustered into an ECAL supercluster. The GSF track is then associated with the supercluster, allowing the tracker and calorimeter information to be combined in the reconstruction of the electron kinematic properties [45, 46]. A schematic representation of the electron reconstruction in the CMS detector is shown in Fig. 2.7.

Electron identification aims to distinguish genuine electrons from background objects, such as hadrons misidentified as electrons and electrons originating from




FIGURE 2.7: Schematic representation of electron passage through the CMS tracker and ECAL, including bremsstrahlung emission and the association of ECAL energy deposits [47].

photon conversions. The identification exploits the compatibility between the reconstructed track and the associated ECAL supercluster, together with properties of the electromagnetic shower and of the surrounding detector activity. Relevant observables include the quality of the GSF track, the geometrical and kinematic matching between the track and the ECAL supercluster, the ratio between the ECAL energy and the track momentum, the number of missing hits in the inner tracker, and variables describing the electromagnetic shower shape. The activity surrounding the electron also provides information to distinguish prompt electrons from electrons produced in hadron decays or embedded in hadronic jets.

In Run 3, CMS provides both cut-based and multivariate electron identification algorithms. Multivariate analysis (MVA) methods combine several discriminating observables into a single discriminator, allowing correlations among the input variables to be exploited. For the electron identification, the MVA discriminator is based on a boosted decision tree (BDT), which is trained to distinguish genuine electrons from background candidates. The input variables include information on the ECAL shower shape, track quality, track–supercluster matching and electron isolation. Different working points (WP) are defined by applying requirements on

the MVA discriminator. They correspond to different compromises between electron identification efficiency and background rejection, with tighter WP providing stronger background suppression at the cost of a lower signal efficiency [48].

Additional criteria are used to reject electrons originating from photon conversions. The conversion rejection exploits the reconstructed track properties and their compatibility with the electron hypothesis, allowing candidates associated with photon conversions to be suppressed.

2.3.3 Complex objects reconstruction

The collection of PF candidates provides a global description of the event and can be used to reconstruct more complex objects, such as jets and missing transverse momentum.

Jets are collimated collections of particles produced by the hadronization of quarks and gluons. Since quarks and gluons cannot be observed as free particles due to colour confinement, their properties are inferred from the resulting hadronic jets [49]. In CMS, jets are reconstructed from PF candidates using the anti- k_T clustering algorithm [50].

The algorithm iteratively combines pairs of particles according to a distance measure defined as

$$d_{ij} = \min(p_{T_i}^{-2}, p_{T_j}^{-2}) \frac{\Delta R_{ij}^2}{R^2}, \quad (2.8)$$

where the angular separation between the two particles in the (η, ϕ) plane is given by

$$\Delta R_{ij} = \sqrt{(\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2}. \quad (2.9)$$

The parameter R defines the characteristic angular size of the reconstructed jets. CMS commonly uses $R = 0.4$ for standard PF jets. The negative power of the transverse momentum in the anti- k_T distance measure causes high p_T particles to be clustered first, resulting in jets with a regular, approximately conical shape. A schematic view of the jet clustering procedure is shown in Fig. 2.8. The four-momentum of the resulting jet is obtained from the momenta of its constituent PF candidates.

Jets originating from the hadronization of bottom quarks, referred to as b jets, can be identified by exploiting the properties of the tracks and secondary vertices associated with the jet. The relatively long lifetime of b hadrons allows them to




FIGURE 2.8: Schematic representation of the anti- k_T clustering procedure. At each step, the distance between pairs of particles is evaluated and the closest pair is combined until the final jets are formed [49].

travel a measurable distance from the primary vertex before decaying, resulting in displaced tracks and secondary vertices that provide a characteristic signature of b jets. Dedicated algorithms combine these observables to distinguish b jets from jets originating from lighter quarks or gluons [51].

Neutrinos and other weakly interacting particles escape the detector without leaving a direct signal. Their presence can nevertheless be inferred from the imbalance of the transverse momentum in the event. Since the initial transverse momentum of the colliding protons is approximately zero, conservation of momentum requires the vector sum of the transverse momenta of all final-state particles to vanish. The missing transverse momentum is consequently defined as

$$\vec{p}_T^{\text{miss}} = - \sum_i \vec{p}_{T_i}, \quad (2.10)$$

where the sum runs over the reconstructed PF candidates. The magnitude of this vector is denoted as the Missing Transverse Energy (MET) [52, 53].

The reconstruction of the MET is particularly sensitive to pile-up, since additional pp interactions contribute to the event momentum balance. In Run 3, the PileUp Per Particle Identification (PUPPI) algorithm is used to mitigate this effect [54].

PUPPI assigns a weight to each reconstructed particle according to the likelihood that it originates from the primary interaction rather than from pile-up, and the particle momenta are rescaled accordingly. The resulting PUPPI candidates are then used to compute the MET. The PUPPI-based reconstruction is the recommended approach for Run 3, as it provides a reduced dependence on pile-up and an improved resolution compared with the standard PF-based reconstruction [46].

2.3.4 Muon reconstruction and identification

Muons are reconstructed using dedicated algorithms that combine information from the inner tracker and the muon system, independently of the PF algorithm. Depending on the detector information used, three types of muons are defined: standalone, tracker and global muons.


Standalone muons are obtained using only information from the muon system. Tracker muons are reconstructed inside-out, extrapolating tracks of the inner tracker towards the muon chambers and requiring a compatible signal in the muon system. Global muons combine the information outside-in, from the muon system to the inner tracker in a global track fit, providing an improved determination of the muon trajectory and momentum.

The different reconstruction approaches are complementary and allow muons to be efficiently reconstructed over a broad kinematic and geometrical range.

After reconstruction, dedicated identification criteria are applied to select muon candidates with different levels of quality. The standard identification working points are loose, medium and tight, corresponding to progressively more stringent requirements on the quality and consistency of the tracker and muon system measurements. These criteria exploit the quality of the reconstructed tracks, the consistency between the inner tracker and muon system information, and the compatibility of the reconstructed trajectory with the interaction region. In particular, the tight identification (tight ID) is designed to select high quality muons and to suppress backgrounds from hadronic particles penetrating the calorimeters and from muons produced in decays in flight [55].

The performance of the muon reconstruction and identification is illustrated in Fig. 2.9. The global reconstruction efficiency is around 98% (Fig. 2.9a), while the tight ID efficiency remains above 90% over most of the detector acceptance

(Fig. 2.9b). The small losses observed in specific pseudorapidity regions are related to detector geometry and to localized inefficiencies in the muon chambers [56].



(a) Global muon reconstruction efficiency as a function of η .

(b) Tight muon identification efficiency as a function of η .

FIGURE 2.9: Muon reconstruction and identification efficiencies as a function of pseudorapidity, measured using the tag-and-probe method in $Z \rightarrow \mu\mu$ events in Run 3 [57].

In addition to the identification criteria, the isolation of the muon is used to distinguish prompt muons from muons produced in the decay of heavy-flavour hadrons or other particles inside jets. Prompt muons produced in the hard interaction are generally isolated from other particles, whereas muons produced within jets are accompanied by additional charged and neutral particles. The relative isolation is therefore defined as the transverse momentum carried by particles in a cone around the muon, normalized to the transverse momentum of the muon:

$$I_{\text{rel}}^{\mu} = \frac{1}{p_T^{\mu}} \left\{ \sum p_T^{\text{charged}} + \max \left[0, \sum p_T^{\text{neutral}} + \sum p_T^{\gamma} - p_T^{PU} \right] \right\}, \quad (2.11)$$

where the sum runs over the reconstructed particles within a cone of radius ΔR around the lepton. Here p_T^{charged} refers to charged hadrons, p_T^{neutral} to neutral hadrons, p_T^{γ} to photons, and p_T^{PU} to the estimated neutral contribution from pile-up [1].

2.4 CMS Run 3 data flow

The large amount of information produced by the CMS detector cannot be stored and processed directly in its raw form. The detector signals are digitized and processed by the CMS reconstruction algorithms, which convert the detector-level information into reconstructed physics objects. The resulting event information is stored in a sequence of data formats with progressively reduced content.

The Analysis Object Data (AOD) format contains a relatively complete description of the reconstructed event. The AOD content is subsequently reduced to produce the MiniAOD format, which provides a more compact representation while retaining the high-level physics objects and the information needed for most physics analyses [58].

For data released from 2016 onwards, CMS introduced the NanoAOD format as a further analysis-oriented data format. NanoAOD is not simply a slimmed down version of MiniAOD, but stores the relevant quantities as branches of ROOT **T**Trees. This makes the format particularly compact and directly accessible with standard ROOT-based analysis tools, without requiring the CMS software environment [59].

The analysis presented in this thesis is performed directly on CMS Run 3 data and simulated samples in the NanoAOD format. The NanoAOD event content provides the reconstructed quantities used to select electrons and muons, reconstruct the dilepton system and apply the event selection described in the following chapter.

Chapter 3

Event selection

This chapter describes the search for the CLFV decay $Z \rightarrow e\mu$ using pp collision data collected by the CMS experiment during the 2024 Run 3 data-taking period. The overall strategy adopted for the analysis is summarised in Fig. 3.1.

The search exploits the characteristic signature of the decay: an opposite-sign electron–muon pair originating from an on-shell Z boson. If present, the signal would be a localised excess in the invariant mass distribution of the $e\mu$ pair, around the Z boson mass. The analysis is designed to isolate this signature while suppressing SM processes that can contribute to the reconstructed $e\mu$ final state, either directly or through misidentification and other reconstruction effects.

The analysis starts from the study of Monte Carlo (MC) samples describing the signal and the main background processes, together with the 2024 collision data. The event selection is developed and performed using the simulated samples and then applied to the data. The selected MC samples are used to evaluate the signal selection efficiency, characterise the background contributions, and optimise the signal-to-background ratio. The signal and background distributions obtained from simulation are then used to construct the invariant mass model for the subsequent statistical analysis, while the selected data are compared with the MC predictions.

The final interpretation is performed with a fit to the invariant mass spectrum, with the aim of either identifying evidence for a $Z \rightarrow e\mu$ signal or, in the absence of a statistically significant excess, setting an upper limit on the branching fraction $\mathcal{B}(Z \rightarrow e\mu)$. The study of the 2024 dataset provides a first analysis framework for this decay using the 109.95 fb^{-1} of data collected in 2024 and represents a step towards its extension to the full Run 3 dataset, corresponding to about 350 fb^{-1} delivered to CMS [21].




FIGURE 3.1: Schematic overview of the analysis strategy adopted for the search for the CLFV decay $Z \rightarrow e\mu$.

3.1 Data and Monte Carlo samples

As discussed in Section 2.4, both data and simulated samples used in this analysis are stored in the NanoAOD format as ROOT n-tuples. The dataset considered in this work corresponds to the 2024 data-taking period of Run 3, at a centre-of-mass energy of 13.6 TeV, with an integrated luminosity of 109.95 fb^{-1} .

Collision data are collected requiring the presence of a high transverse momentum electron or muon, through the single lepton triggers `HLT_Ele30_WPTight_Gsf` and `HLT_IsoMu24`. The two triggers define two separate primary datasets, an electron- and a muon-triggered one, which are combined in the analysis, avoiding double counting, to select events containing an electron–muon pair. The full trigger logic, including its consistent application to data and to the simulated samples, is discussed in Section 3.2.1. The data samples used are summarised in Table 3.1.

TABLE 3.1: Data samples used in the analysis. The integrated luminosity corresponding to the 2024 data-taking period is reported.

Dataset	Year	$L_{\text{int}} [\text{fb}^{-1}]$
Data_EGamma_2024	2024	109.95
Data_Muon_2024	2024	109.95

MC simulated samples are used to develop and optimise the event selection, evaluate the signal efficiency, model the electron–muon pair invariant mass distribution ($m_{e\mu}$), and characterise the background processes. The signal and background samples are described separately in Sections 3.1.1 and 3.1.2.

To compare the simulated samples with data, the MC events are normalised to the integrated luminosity of the data sample. This normalisation accounts for the production rate of the corresponding process in pp collisions and for the decay channel considered, while also taking into account the generator weights and the corrections applied to the simulated events.

The weight assigned to each simulated event is therefore defined as

$$w = \frac{w_{\text{gen}}}{\sum w_{\text{gen}}} \sigma_{\text{pr}} \mathcal{B} L_{\text{int}} \text{SF} w_{\text{PU}}, \quad (3.1)$$

where w_{gen} is the generator weight, i.e. the weight assigned to each event by the MC generator, which accounts for the relative contribution of that event to the simulated process (for next-to-leading-order generators it can also take negative values), and $\sum w_{\text{gen}}$ is the sum of generator weights over the whole sample, which

normalises the sample to the expected number of events. The normalisation is fixed by σ_{pr} the production cross section, the branching fraction $\mathcal{B}$ of the considered process decay, and the integrated luminosity L_{int} of the data sample. The factor SF accounts for corrections to the simulated lepton identification and isolation efficiencies and is given by

$$\text{SF} = \text{SF}_e^{\text{ID}} \text{SF}_\mu^{\text{iso}} \text{SF}_\mu^{\text{tightID}},$$

where the three factors correct the electron identification, muon isolation, and muon identification efficiencies, respectively. Finally, w_{PU} is the pile-up weight. The simulated samples are generated with an assumed distribution of the number of pp interactions per bunch crossing, which does not coincide with the one observed in data. Each simulated event is therefore reweighted by the ratio between the pile-up distribution measured in data and the simulated one, evaluated at the number of interactions of that event.

3.1.1 Signal sample

The signal process considered in this analysis is the CLFV decay $Z \rightarrow e\mu$. Since this decay is not expected to occur at an observable rate in the SM, a dedicated MC sample is used to model the signal kinematics and evaluate the selection efficiency. The signal sample contains 30 000 generated events and is processed through the standard CMS simulation and reconstruction chain. The branching fraction of the CLFV decay is the quantity to be determined by the analysis; a reference value of $\mathcal{B}(Z \rightarrow e\mu) = 10^{-6}$ is therefore adopted to normalise the signal sample. This value is used as the branching fraction in the event weight of Eq. (3.1).

The main characteristics of the signal and background samples are summarised in Table 3.2.

3.1.2 Background samples

The background processes considered in this analysis are SM processes that can contribute to the selected electron–muon final state. The simulated samples include Drell–Yan production with $Z \rightarrow ee$, $Z \rightarrow \mu\mu$, and $Z \rightarrow \tau\tau$ decays, diboson production through $WW \rightarrow 2\ell 2\nu$ and $WZ \rightarrow 3\ell\nu$, and top-quark pair production through $t\bar{t} \rightarrow 2\ell 2\nu 2b$.

TABLE 3.2: MC samples used in the analysis. The process, production cross section, branching fraction, and number of generated events are reported [19, 20, 60, 61, 62].

Sample	Process	σ_{pr} [pb]	$\mathcal{B}$ [%]	N_{gen}
ZEMuAnalysis_Signal	$pp \rightarrow Z \rightarrow e\mu$	6×10^4	10^{-4}	30 000
DY_JetsToEE_2024	$pp \rightarrow Z \rightarrow ee$	6×10^4	3.36	477 172 288
DY_JetsToMuMu_2024	$pp \rightarrow Z \rightarrow \mu\mu$	6×10^4	3.37	492 270 708
DY_JetsToTauTau_2024	$pp \rightarrow Z \rightarrow \tau\tau$	6×10^4	3.37	486 994 937
TTbar_2L2Nu_2024	$pp \rightarrow t\bar{t} \rightarrow 2\ell 2\nu$	924	10.61	470 123 263
WWTo2L2Nu_2024	$pp \rightarrow WW \rightarrow 2\ell 2\nu$	126	10.61	239 943 886
WZTo3LNu_2024	$pp \rightarrow WZ \rightarrow 3\ell\nu$	55	3.29	248 149 069

The $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ samples can enter the selected sample when a lepton is misidentified or when additional reconstructed leptons result in an electron–muon pair passing the selection. In the $Z \rightarrow \tau\tau$ sample, a genuine $e\mu$ pair can arise from the leptonic decays of the two τ leptons.

The $WW \rightarrow 2\ell 2\nu$ sample describes events in which both W bosons decay leptonically. An $e\mu$ pair can be produced directly when the two W bosons decay into an electron and a muon, respectively, or through intermediate τ decays.

The $WZ \rightarrow 3\ell\nu$ sample can contribute when the W and Z bosons decay leptonically and one of the three charged leptons is not reconstructed or fails the selection, leaving an electron–muon pair. Contributions involving intermediate τ decays are also possible.

The $t\bar{t} \rightarrow 2\ell 2\nu 2b$ sample describes top-quark pair events in which both top quarks decay through $t \rightarrow Wb$ and both W bosons decay leptonically. These events can contain an $e\mu$ pair together with the two b quarks. The presence of additional hadronic activity makes this process particularly relevant for the event selection. For the WW , WZ , and $t\bar{t}$ samples, the leptonic final states are generated inclusively in lepton flavour. The corresponding inclusive branching fractions are therefore used for the MC normalisation.

For $WW \rightarrow 2\ell 2\nu$, and analogously for $t\bar{t} \rightarrow WbW\bar{b} \rightarrow 2\ell 2\nu b\bar{b}$, the inclusive branching fraction is

$$\mathcal{B}(WW \rightarrow 2\ell 2\nu) = 9 \mathcal{B}(W \rightarrow \ell\nu)^2,$$

with $\mathcal{B}(W \rightarrow \ell\nu) \simeq 10.86\%$ [19]. The factor 9 accounts for the 3×3 possible combinations of electron, muon, and tau flavours from the two W bosons. The same

leptonic branching fraction applies to the $t\bar{t}$ sample since $\mathcal{B}(t \rightarrow Wb) \simeq 100\%$. For $WZ \rightarrow 3\ell\nu$, the inclusive branching fraction is

$$\mathcal{B}(WZ \rightarrow 3\ell\nu) = 9 \mathcal{B}(W \rightarrow \ell\nu) \mathcal{B}(Z \rightarrow \ell\ell),$$

with $\mathcal{B}(Z \rightarrow \ell\ell) \simeq 3.37\%$ [19]. These inclusive branching fractions are the values used in the event weights and reported in Table 3.2, together with the branching fractions for the Drell–Yan samples.

For illustration, it is useful to quantify the fraction of events that can directly produce the $e\mu$ topology.

For $Z \rightarrow \tau\tau$, using $\mathcal{B}(\tau \rightarrow e\nu_\tau\bar{\nu}_e) = 17.82\%$ and $\mathcal{B}(\tau \rightarrow \mu\nu_\tau\bar{\nu}_\mu) = 17.39\%$ [19],

$$\mathcal{B}(\tau\tau \rightarrow e\mu 4\nu) = 2 \mathcal{B}(\tau \rightarrow e\nu_\tau\bar{\nu}_e) \mathcal{B}(\tau \rightarrow \mu\nu_\tau\bar{\nu}_\mu) \simeq 6.20\%.$$

For WW , the leptonic branching fractions $\mathcal{B}(W \rightarrow e\nu_e) \simeq 10.71\%$ and $\mathcal{B}(W \rightarrow \mu\nu_\mu) \simeq 10.63\%$ [19] give a direct $e\mu$ fraction of

$$\mathcal{B}(WW \rightarrow e\mu 2\nu) = 2 \mathcal{B}(W \rightarrow e\nu_e) \mathcal{B}(W \rightarrow \mu\nu_\mu) \simeq 2.28\%,$$

with additional, smaller contributions from channels involving one or two intermediate τ leptons. The same direct $e\mu$ fraction applies to the $t\bar{t}$ sample, giving approximately 2.29%.

For $WZ \rightarrow 3\ell\nu$, a direct $e\mu$ pair arises from $W \rightarrow e\nu_e$ combined with $Z \rightarrow \mu\mu$, or $W \rightarrow \mu\nu_\mu$ combined with $Z \rightarrow ee$. In both cases, the remaining lepton from the Z decay must not be reconstructed or must fail the selection. Using $\mathcal{B}(Z \rightarrow \ell\ell) \simeq 3.37\%$ [19],

$$\mathcal{B}(WZ \rightarrow e\mu\ell\nu) = \mathcal{B}(W \rightarrow e\nu_e) \mathcal{B}(Z \rightarrow \mu\mu) + \mathcal{B}(W \rightarrow \mu\nu_\mu) \mathcal{B}(Z \rightarrow ee) \simeq 0.72\%.$$

Also in this case, additional subdominant contributions arise from channels involving one or more intermediate τ decays.

In all final states containing neutrinos, a larger value of missing energy in the collision is expected.

3.2 Selection algorithm

The event selection is designed to identify events containing an electron–muon pair compatible with the $Z \rightarrow e\mu$ decay while suppressing the background contributions. The selection is applied in four steps: trigger selection, lepton selection, event selection, and pair selection. The selection criteria are first optimised using MC simulation and are subsequently applied to collision data.

The same trigger requirement is applied consistently to collision data and MC simulation, with an implementation adapted to the structure of the corresponding samples, as discussed in Section 3.2.1.

3.2.1 Trigger selection

The first step of the event selection is the trigger requirement. Events are required to pass at least one of the two single-lepton triggers used in this analysis. The single-electron trigger requires an electron candidate with $p_T > 30$ GeV, whose track is reconstructed with the GSF algorithm (Section 2.3.2) and which satisfies the tight working point of the HLT electron selection, a set of requirements on the shower shape, the track–cluster matching and the isolation of the candidate. The single-muon trigger requires a muon candidate with transverse momentum above 24 GeV that is isolated, i.e. with limited activity in the tracker and in the calorimeters in a cone around its direction (Section 2.3.4).

In simulated events, which are generated as a single inclusive sample, this requirement is implemented directly as the logical OR

$$\text{HLT_Ele30_WPTight_Gsf} \vee \text{HLT_IsoMu24}.$$

Collision data, on the other hand, are recorded in two separate primary datasets, corresponding to the electron and muon triggers (Table 3.1). To combine the two datasets without double counting the common events, events in the electron-triggered dataset are required to satisfy `HLT_Ele30_WPTight_Gsf` while explicitly failing `HLT_IsoMu24`; events in the muon-triggered dataset are simply required to satisfy `HLT_IsoMu24`.

3.2.2 Lepton selection

The $Z \rightarrow e\mu$ signal is expected to produce one prompt electron and one prompt muon of opposite charge, both originating from the primary interaction vertex. The lepton selection is therefore designed to identify well-reconstructed, isolated leptons within the detector acceptance while suppressing backgrounds from misidentified or non-prompt leptons. The requirements applied to electron and muon candidates are summarized in Tables 3.3 and 3.4, respectively. These requirements follow the selection strategy of the previous CMS search for $Z \rightarrow e\mu$ with Run 2 data [1], and are adapted to the 2024 Run 3 dataset and to the available NanoAOD variables. The corresponding variables are those introduced in Chapter 2. In particular, the selection exploits the lepton kinematics, geometrical acceptance, identification and isolation information, as well as the compatibility of the reconstructed lepton tracks with the primary interaction vertex.

TABLE 3.3: Selection criteria applied to electron candidates.

Variable	Requirement
$p_T(e)$	$> 30 \text{ GeV}$
$ \eta(e) $	< 2.5
$ \eta_{SC}(e) $	$\notin [1.444, 1.566]$
Conversion veto	pass
MVA Iso ID	WP90
$ d_{xy} $	$< 0.2 \text{ cm}$
$ d_z $	$< 0.5 \text{ cm}$
$ d_{xy}/\sigma_{d_{xy}} $	< 3
$ d_z/\sigma_{d_z} $	< 5

TABLE 3.4: Selection criteria applied to muon candidates.

Variable	Requirement
$p_T(\mu)$	$> 30 \text{ GeV}$
$ \eta(\mu) $	< 2.4
Muon ID	Tight
Relative PF Iso ($\Delta R = 0.3$)	< 0.05
$ d_{xy} $	$< 0.2 \text{ cm}$
$ d_z $	$< 0.5 \text{ cm}$
$ d_{xy}/\sigma_{d_{xy}} $	< 3
$ d_z/\sigma_{d_z} $	< 5

Electron candidates are first required to be within the tracker and ECAL geometric acceptance, with $|\eta| < 2.5$. The transition region between the ECAL barrel and endcaps is excluded by requiring $|\eta_{\text{SC}}| \notin [1.444, 1.566]$. This region is characterized by a change in detector geometry and a reduced reconstruction performance, and is therefore not used in the analysis.

The selected electrons are required to pass a conversion veto. Photons can convert into an e^+e^- pair when interacting with the material of the detector. The two charged particles produced in the conversion leave tracks in the tracker, whose properties can differ from those expected for a prompt electron produced at the primary interaction point. The conversion veto exploits information from the reconstructed electron track and its compatibility with the electron hypothesis to identify and reject candidates consistent with photon conversions. In this way, electrons originating from $\gamma \rightarrow e^+e^-$ processes are suppressed.

Electron candidates are also required to pass an MVA-based isolation identification. As described in Section 2.3.2, the MVA identification combines several observables related to the electron track, the ECAL shower, the track–supercluster matching and the activity around the electron into a single discriminator. The WP90 working point is chosen to retain a high efficiency for genuine electrons while rejecting background candidates.

Finally, requirements are imposed on the transverse and longitudinal impact parameters, d_{xy} and d_z , and on their corresponding significances, $\sigma_{d_{xy}}$ and σ_{d_z} . The impact parameters quantify the displacement of the reconstructed lepton trajectory with respect to the primary interaction vertex, while the significance accounts for the uncertainty of their measurement. Requiring both the absolute values and their significances to be small selects electrons whose tracks are compatible with originating from the primary interaction point. This suppresses electrons produced at displaced vertices, for example in the decay of long-lived particles or in hadronic activity.

Muon candidates are required to satisfy the geometrical acceptance of the CMS muon system, with $|\eta| < 2.4$, and to pass a tight muon identification requirement, defined in Section 2.3.4. The identification criteria combine information from the inner tracker and the muon system to select muons with a well-reconstructed trajectory and consistent measurements in the different detector components. The

tight working point provides stronger rejection of incorrectly reconstructed or non-prompt muons.

To further suppress non-prompt muons, candidates are required to satisfy a relative PF isolation requirement. The isolation variable, defined in Eq. (2.11), quantifies the amount of reconstructed activity surrounding the muon within a cone of $\Delta R = 0.3$, relative to the muon transverse momentum. Prompt muons produced in the decay of a Z boson are typically isolated from other particles, whereas muons produced in hadron decays inside jets are accompanied by additional hadronic activity. The requirement on the relative isolation therefore preferentially rejects muons originating from jets while retaining prompt muons from the signal. This choice is motivated by the relative isolation distribution shown in Fig. 3.2, obtained before applying the lepton selection. Signal muons, being prompt, populate the low-isolation region, while the dominant background processes, $Z \rightarrow ee$, $Z \rightarrow \tau\tau$, and $t\bar{t} \rightarrow 2\ell 2\nu 2b$, extend towards larger isolation values. In $Z \rightarrow ee$ the muon candidate is non-prompt, in $Z \rightarrow \tau\tau$ it originates from the three-body decay $\tau \rightarrow \mu\nu\nu$, and in $t\bar{t} \rightarrow 2\ell 2\nu 2b$ the overall hadronic activity is enhanced by the presence of two b -quark jets in the event, increasing the probability of residual activity within the isolation cone. The dashed line in Fig. 3.2 indicates the chosen threshold of 0.05.

As for electrons, the compatibility of the muon track with the primary interaction vertex is further enforced through requirements on d_{xy} , d_z , and their corresponding significances. These requirements suppress muons originating from displaced decays or from secondary interactions, while retaining muons produced promptly at the primary vertex.

The common requirement $p_T > 30 \text{ GeV}$ is motivated by the transverse momentum distributions shown in Fig. 3.3, obtained before applying the lepton selection. Since the $Z \rightarrow e\mu$ signal originates from the two-body decay of a Z boson, the two leptons are expected to have relatively high transverse momenta, of the order of $\frac{m_Z}{2}$ in the Z rest frame. The electron and muon p_T distributions are shown separately in Figs. 3.3a and 3.3b. In each distribution, the signal is compared with the background processes most directly affected by the p_T requirement.

For the electron p_T distribution, the $Z \rightarrow \mu\mu$ and $Z \rightarrow \tau\tau$ backgrounds are considered. The $Z \rightarrow \mu\mu$ process does not contain electrons in the final state. It can enter the $e\mu$ selection when a muon is misidentified or incorrectly reconstructed




FIGURE 3.2: Relative isolation distribution of reconstructed muons before applying the lepton selection. The distribution is shown for the signal and the background processes most relevant to the muon isolation requirement: $Z \rightarrow ee$, $Z \rightarrow \tau\tau$, and $t\bar{t} \rightarrow 2\ell 2\nu$. The dashed line indicates the chosen threshold of isolation < 0.05 .

as an electron. These reconstructed electron candidates are predominantly found at low p_T , resulting in a distribution that falls rapidly above the low p_T region. The $Z \rightarrow \tau\tau$ background, instead, can produce a genuine electron through the leptonic decay $\tau \rightarrow e\nu\bar{\nu}$. The neutrinos carry away part of the energy and momentum of the τ lepton, resulting in a softer electron p_T spectrum than that expected for signal.

A similar behaviour is observed for the muon p_T distribution. The $Z \rightarrow ee$ background can contribute when an electron is misidentified or incorrectly reconstructed as a muon. As shown in Fig. 3.3b, the corresponding reconstructed muons are predominantly concentrated at low p_T . In $Z \rightarrow \tau\tau$ events, genuine muons can be produced through the decay $\tau \rightarrow \mu\nu\bar{\nu}$. Again, the presence of neutrinos reduces the momentum carried by the visible muon, producing a softer p_T spectrum than that of the signal.

The $p_T > 30 \text{ GeV}$ requirement therefore removes a substantial fraction of the low p_T background while retaining a large fraction of signal. It provides discrimination

against both backgrounds containing misidentified or incorrectly reconstructed leptons and the softer leptons produced in τ decays. It is also consistent with the transverse momentum thresholds of the single-lepton triggers used in the analysis.

The complete set of electron and muon requirements defines the selected lepton candidates used in the subsequent event selection.


3.2.3 Event-level selection

After the lepton selection, additional event-level requirements are applied to suppress background processes whose topology differs from that expected for the $Z \rightarrow e\mu$ signal. The corresponding requirements are summarised in Table 3.5.


TABLE 3.5: Event-level selection criteria.

Variable	Requirement
Jet- p_T veto	No jet with $p_T^{\text{jet}} > 70$ GeV
b -jet veto	$p_T^{\text{jet}} > 15$ GeV for identified b jets
E_T^{miss}	< 40 GeV
Number of leptons	Exactly one selected electron and exactly one selected muon

The first requirement is a veto on events containing a jet with transverse momentum above 70 GeV. This selection is introduced to suppress the $t\bar{t}$ background. In $t\bar{t}$ events, the top quarks decay predominantly as $t \rightarrow Wb$, producing two b quarks in addition to the two W bosons. Since the top quark has a mass of about 173 GeV, its decay produces energetic b jets. When both W bosons decay leptonically, the final state can therefore contain an electron–muon pair accompanied by b jets. As a consequence, the b jets from top-quark decays tend to have a harder transverse momentum spectrum than the jets expected in the $Z \rightarrow e\mu$ signal. Therefore, events containing a jet with $p_T > 70$ GeV are rejected, reducing the contribution from $t\bar{t}$ events while retaining most of the signal. The distribution of the transverse momentum of the leading jet, defined as the jet with the highest p_T in each event, before the jet- p_T veto is shown in Fig. 3.4a for the signal and for the $t\bar{t}$ background. The signal is concentrated at lower p_T values, mostly below the threshold, while the $t\bar{t}$ distribution extends to higher values. The requirement therefore removes a large fraction of the background while retaining most of the signal events.



(a) Electron p_T distribution for the $Z \rightarrow e\mu$ signal, $Z \rightarrow \mu\mu$, and $Z \rightarrow \tau\tau$ backgrounds.



(b) Muon p_T distribution for the $Z \rightarrow e\mu$ signal, $Z \rightarrow ee$, and $Z \rightarrow \tau\tau$ backgrounds.

FIGURE 3.3: Transverse momentum distributions of reconstructed electrons and muons before applying the lepton selection. The distributions are shown for the signal and the background processes most relevant to the corresponding lepton p_T requirement. The dashed line indicates the common threshold of $p_T > 30$ GeV.

The second requirement is a veto on events containing b -tagged jets. The jets considered in this analysis are AK4 jets, reconstructed with the anti- k_T algorithm with a distance parameter $R = 0.4$, using PF candidates as inputs. The effect of pile-up is mitigated using the PUPPI algorithm, which assigns a weight to each particle according to its probability of originating from a pile-up interaction or from the leading vertex, as described in Section 2.3.3.

The b -jet identification, known as b -tagging, is performed using the UParT algorithm. UParT is a machine learning model designed for AK4 jets. It performs jet flavour identification exploiting information from the jet constituents and secondary vertices [63].

I consider a jet to be b -tagged when it satisfies the two requirements

$$p_T^{\text{jet}} > 15 \text{ GeV} \quad \text{and} \quad b\text{-tag score} > 0.2.$$

The transverse momentum threshold restricts the b -tagging selection to jets with $p_T > 15 \text{ GeV}$, where the jet reconstruction and flavour-tagging performance is well controlled. The threshold on the UParT discriminator score > 0.2 corresponds to a loose b -tagging working point. Using this working point in a veto provides a high efficiency for identifying b jets and therefore allows a large fraction of events containing b jets to be rejected.

This requirement is applied as a veto: an event is rejected if at least one jet satisfies both conditions. The b -jet veto therefore provides an additional suppression of the $t\bar{t}$ background, exploiting the presence of b jets in these events. Since the $Z \rightarrow e\mu$ signal does not contain b quarks in the decay of the Z boson at leading order, this veto is expected to have a limited effect on the signal acceptance while efficiently rejecting $t\bar{t}$ events.

Figure 3.4b shows the number of events with zero or at least one b -tagged jet for the signal and the $t\bar{t} \rightarrow 2\ell 2\nu 2b$ background before applying the veto. The signal is found to be concentrated almost entirely in the zero b -jet bin, while the $t\bar{t} \rightarrow 2\ell 2\nu 2b$ background is predominantly characterised by the presence of at least one b -tagged jet. This demonstrates the discriminating power of the b -jet veto against the $t\bar{t}$ background.

Events with $E_T^{\text{miss}} > 40 \text{ GeV}$ are rejected. The missing transverse energy is computed using the PUPPI algorithm.

This requirement suppresses backgrounds containing neutrinos in the final state, such as $WW \rightarrow 2\ell 2\nu$ and $t\bar{t} \rightarrow 2\ell 2\nu 2b$, which populate the high E_T^{miss} region. The $Z \rightarrow e\mu$ signal, instead, is concentrated at low E_T^{miss} , apart from contributions due to detector resolution. The $Z \rightarrow \tau\tau \rightarrow e\mu + 4\nu$ background has an intermediate E_T^{miss} distribution, since the neutrinos produced in the τ decays carry away part of the momentum. Consequently, this background is reduced by the requirement but is not completely suppressed. The E_T^{miss} distribution before applying the requirement is shown in Figure 3.4c for the signal and the relevant background processes.

Finally, events containing more than one selected electron or more than one selected muon are rejected. This ensures that the event contains exactly one selected electron and one selected muon, removing ambiguities before constructing the $e\mu$ pair.

3.2.4 Pair selection

Following the lepton and event selections, the electron and muon candidates are combined to form an $e\mu$ pair. Since events with more than one selected electron or more than one selected muon are rejected, at most one electron–muon pair can be formed in each event.

The pair is first required to have opposite electric charge. This requirement is implemented by imposing


$$q_e q_\mu < 0,$$

where q_e and q_μ are the charges of the electron and muon, respectively. This selects the opposite sign (OS) topology expected for the $Z \rightarrow e\mu$ decay and suppresses combinations of leptons with the same charge.


The invariant mass of the electron–muon pair is then required to lie in the range

$$70 \text{ GeV} < m_{e\mu} < 110 \text{ GeV},$$


following the same mass window adopted in the corresponding Run 2 analysis [1].



(a) Leading jet p_T distribution before the jet- p_T veto, where the leading jet is defined as the jet with the highest p_T in each event, for the $Z \rightarrow e\mu$ signal and $t\bar{t}$ background. The dashed vertical line indicates the 70 GeV threshold for selection.



(b) Number of events with zero or at least one b -tagged jet before applying the b jet veto, for the $Z \rightarrow e\mu$ signal and $t\bar{t} \rightarrow 2\ell 2\nu 2b$.



(c) Missing transverse energy distribution before the $E_T^{\text{miss}} < 40$ GeV requirement, for signal and the relevant backgrounds. The dashed vertical line indicates the cut threshold.

FIGURE 3.4: Distributions of the variables used in the event-level selection, shown before applying the corresponding requirements.

The invariant mass is calculated from the transverse momenta, pseudorapidities, and azimuthal angles of the two leptons as

$$m_{e\mu} = \sqrt{2p_T^e p_T^\mu [\cosh(\eta_e - \eta_\mu) - \cos(\phi_e - \phi_\mu)]}. \quad (3.2)$$

The selected mass window is centred on the Z boson mass and defines the region used in the subsequent analysis of the invariant mass distribution of the $e\mu$ pair.

After the complete event and pair selection, the invariant mass of the electron–muon system provides the main observable used to identify the $Z \rightarrow e\mu$ signal. Since both final-state leptons are directly reconstructed, the four-momentum of the $e\mu$ system can be determined and its invariant mass can be calculated without relying on the reconstruction of invisible particles. Signal events therefore produce a localized peak around the Z -boson mass, while the background contributions give broader and smoothly varying distributions. In particular, the $Z \rightarrow \tau\tau$ contribution decreases rapidly with increasing $m_{e\mu}$, while processes such as $WW \rightarrow 2\ell 2\nu$

contribute over a broader mass range. The $m_{e\mu}$ distribution obtained from the simulated samples after the complete selection is shown in Fig. 3.5.




FIGURE 3.5: Stacked distribution of the invariant mass $m_{e\mu}$ after the complete event and pair selection for the simulated signal and background samples. The signal contribution is normalized to the reference branching fraction $\mathcal{B}(Z \rightarrow e\mu) = 10^{-6}$.

The effectiveness of the final selection can also be assessed through an estimate of the expected statistical significance of a possible signal excess over the background. The significance provides an indication of how well a potential signal can be distinguished from the background. Under the background-only hypothesis, it quantifies the compatibility of an observed excess with a statistical fluctuation of the background.

In a simplified approximation, the expected significance can be estimated as

$$Z = \frac{S}{\sqrt{B}}, \quad (3.3)$$

where S and B denote the expected numbers of signal and background events, respectively, after the complete selection. The signal and background yields are obtained from the corresponding simulated samples after applying the event weights described in Eq. (3.1), and are therefore normalised to the integrated luminosity considered in this analysis. The resulting value provides an approximate measure of the separation between the signal and the background.

In particle physics, a significance of approximately 3σ is conventionally associated with evidence for a phenomenon, while a significance of 5σ is conventionally required to claim its discovery.

For the complete selection and for the signal's reference branching fraction $\mathcal{B}(Z \rightarrow e\mu) = 10^{-6}$, the resulting approximate significance is $Z = 6.84\sigma$. This estimate accounts only for the statistical fluctuation of the background and does not include systematic uncertainties.

3.2.5 Selection efficiency

To evaluate the effect of the different selection requirements, I record the number of surviving events after each step of the selection, separately for signal and for each simulated background sample. The reference is the number of generated events, N_{gen} , reported in Table 3.2. Only a fraction of the generated events is available in the analysed samples, since a pre-selection is applied during the sample production. This fraction, referred to as pre-selection efficiency, is defined as the ratio between the number of events available in the analysed sample and the total number of generated events, and it reflects the detector acceptance and the trigger requirement. Since the trigger is part of the pre-selection, all the events available to the analysis already satisfy it, and it is not reported as a separate selection step. The efficiency of each subsequent selection step, relative to the previous one, is defined as

$$\varepsilon_i^{\text{rel}} = \frac{N_i}{N_{i-1}}, \quad (3.4)$$

where N_i and N_{i-1} are the numbers of events surviving the current and the preceding selection requirements, respectively. The total selection efficiency is defined as

$$\varepsilon_{\text{tot}} = \frac{N_{\text{final}}}{N_{\text{gen}}}, \quad (3.5)$$

where N_{final} is the number of events surviving the complete selection chain. This quantity measures the overall fraction of generated events retained after both the pre-selection and the analysis selection. These efficiencies are computed from the unweighted numbers of simulated events.

The results are summarised in Table 3.6, which reports the pre-selection efficiency, the relative efficiency of each selection step, as defined in Eq. (3.4), the final number of selected events and the total selection efficiency, as defined in Eq. (3.5). This allows the impact of each requirement to be compared across the different processes.

Table 3.6 highlights several features of the selection. The $t\bar{t} \rightarrow 2\ell 2\nu 2b$ background shows the largest reduction at the jet- p_T veto step, with a relative efficiency of only 25.7%. The b -jet veto also provides a significant rejection, with a relative efficiency of 48.9%. Both effects are consistent with the presence of two b quarks in top quark decays, as discussed previously. A further strong reduction follows from the E_T^{miss} requirement, with a relative efficiency of 31.9%, due to the neutrinos produced in the leptonic W decays.

The $WZ \rightarrow 3\ell\nu$ sample is strongly reduced by several event-level requirements. In particular, the jet- p_T veto and the E_T^{miss} requirement have relative efficiencies of 42.7% and 45.3%, respectively. A further substantial reduction occurs when requiring exactly one electron and one muon, with a relative efficiency of 66.4%, consistent with the presence of a third lepton.

The two purely leptonic Drell–Yan backgrounds, $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$, have the smallest efficiency at the lepton selection step, 0.4% and 10.6%, respectively. An $e\mu$ pair in these samples can arise from lepton misidentification, misreconstruction, or additional reconstructed leptons, and is therefore strongly suppressed by the lepton selection. The subsequent opposite-sign requirement provides an additional strong rejection, with relative efficiencies of 52.2% for $Z \rightarrow ee$ and 60.1% for $Z \rightarrow \mu\mu$.

The $Z \rightarrow \tau\tau$ background instead produces genuine electrons and muons through leptonic τ decays, resulting in a relatively higher lepton selection efficiency of 14.8%. However, the neutrinos produced in the τ decays carry away part of the momentum, leading to a broader $e\mu$ invariant mass distribution and to a lower efficiency of the mass-window requirement, which is 64.4%.

The signal has the largest total selection efficiency, $\varepsilon_{\text{tot}} \simeq 15.8\%$. This is expected from the presence of two genuine leptons produced directly in the two-body $Z \rightarrow e\mu$ decay, with kinematics well suited to the selection requirements.

TABLE 3.6: Pre-selection efficiency, relative efficiency of each selection step with respect to the preceding one, final number of selected events and total selection efficiency, as defined in Eq. (3.5), for the signal and background samples. The pre-selection efficiency is calculated with respect to N_{gen} , reported in Table 3.2.

Sample	Pre-sel.	Lepton sel.	Jet- p_T veto	b -jet veto	E_T^{miss}	Single e, μ	OS pair	Mass window	N_{final}	ε_{tot}
$Z \rightarrow e\mu$	37.5%	60.9%	86.1%	98.0%	82.7%	100%	99.75%	99.0%	4724	1.6×10^{-1}
$Z \rightarrow ee$	0.03%	0.4%	74.1%	72.1%	82.1%	96.8%	52.2%	85.9%	110	2.3×10^{-7}
$Z \rightarrow \mu\mu$	0.02%	10.6%	79.4%	86.1%	81.1%	98.4%	60.1%	94.0%	3217	6.5×10^{-6}
$Z \rightarrow \tau\tau$	0.1%	14.8%	91.4%	95.8%	80.1%	100%	99.7%	64.4%	24580	5.1×10^{-5}
$t\bar{t} \rightarrow 2\ell 2\nu 2b$	0.9%	38.4%	25.7%	48.9%	31.9%	99.9%	99.6%	74.9%	51506	1.1×10^{-4}
$WW \rightarrow 2\ell 2\nu$	4.0%	43.5%	62.2%	98.2%	49.8%	100%	99.8%	78.5%	1002444	4.2×10^{-3}
$WZ \rightarrow 3\ell\nu$	1.9%	39.1%	42.7%	93.3%	45.3%	66.4%	49.3%	74.1%	83700	3.4×10^{-4}

3.2.6 MC and Data comparison

The complete event selection described in the previous sections is applied both to the simulated samples and to the collision data collected in 2024. The selected data are compared with the prediction obtained from the simulated background processes to validate the selection and the background modelling before the statistical analysis.

The main comparison is performed using the electron–muon invariant mass which is the final discriminating variable of the analysis. Figure 3.6 shows the resulting distribution after the complete selection, with the simulated background contributions compared with the observed data. The MC samples are normalised to the integrated luminosity of the data sample as described in Eq. (3.1). A fair agreement between data and simulation is observed.

To avoid a bias in the development and validation of the statistical model, the region $80 \text{ GeV} < m_{e\mu} < 100 \text{ GeV}$ is kept blind at this stage of the analysis. The comparison outside this region provides a check of the background prediction independently of the events that may contain a signal contribution.

The validated $m_{e\mu}$ distribution is then used as the basis for the statistical analysis. The signal and background models, together with the fitting procedure, are introduced in Chapter 4, followed by the statistical interpretation and the resulting upper limit on $\mathcal{B}(Z \rightarrow e\mu)$ in Chapter 5.




FIGURE 3.6: Invariant mass of the $e\mu$ system for data (points, with bars denoting statistical uncertainty) and simulated background (stacked filled histograms) after the complete event selection. The region $80 < m_{e\mu} < 100$ GeV, where a potential $Z \rightarrow e\mu$ signal is expected, is blinded and excluded from the data–MC comparison. The lower panel shows the ratio of the data to the total simulated background yield in each bin.

Chapter 4

Signal and background modelling

The search for the CLFV decay $Z \rightarrow e\mu$ exploits the invariant mass distribution of the selected electron–muon pairs, which is the final discriminating observable of the analysis.

As discussed in Chapter 3, the event selection is designed to isolate the characteristic $e\mu$ topology of the signal while suppressing the main SM background contributions. The resulting $m_{e\mu}$ distribution is therefore used to distinguish a possible signal contribution from the background and forms the basis of the statistical analysis.

Figure 3.5 of Chapter 3 shows that the simulated $Z \rightarrow e\mu$ signal produces a narrow peak at the Z boson mass, $m_Z \approx 91.2$ GeV, since the two selected leptons originate directly from the decay of an on-shell Z boson. The combined background, on the other hand, gives a smooth distribution. This difference in shape allows a possible $Z \rightarrow e\mu$ contribution to be extracted through a parametric fit to the $m_{e\mu}$ distribution.

The signal and background models are determined from simulation, combined into a single model, validated on simulated samples and pseudo-experiments, and finally applied to data.

4.1 Fitting strategy

The statistical modelling is performed using the RooFit framework within ROOT. RooFit provides the tools required to define probability density functions (PDF) and determine their parameters through maximum-likelihood fits [64]. The observable used in the fit is the invariant mass of the selected electron–muon pair, i.e., $m_{e\mu}$.

As implemented in RooFit, each PDF $f(m_{e\mu}; \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ denotes the set of model parameters, is normalised over the fit range $[m_{\min}, m_{\max}]$,

$$\int_{m_{\min}}^{m_{\max}} f(m_{e\mu}; \boldsymbol{\theta}) \, dm_{e\mu} = 1, \quad (4.1)$$

and can therefore be interpreted as the probability density for observing an event at a given value of $m_{e\mu}$.

The parameters are determined by an unbinned maximum-likelihood fit. For a sample of N events the likelihood is

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{i=1}^N f(m_{e\mu,i}; \boldsymbol{\theta}), \quad (4.2)$$

and the best-fit values are obtained by minimising $-\ln \mathcal{L}$. The parameter uncertainties are estimated from the second derivatives of $-\ln \mathcal{L}$ at the minimum and are statistical only. The fit status and the quality of the covariance matrix are checked to assess the reliability of the fit results.

This procedure is first used to determine the signal and background shapes separately, from the corresponding simulated samples (Sections 4.2 and 4.3).

Since the samples are simulated, each event enters the likelihood with a weight w_i , which accounts for the normalisation of the sample to the integrated luminosity of the data and for the correction factors described in Eq. (3.1). The negative log-likelihood is therefore evaluated as

$$-\ln \mathcal{L}(\boldsymbol{\theta}) = -\sum_{i=1}^N w_i \ln f(m_{e\mu,i}; \boldsymbol{\theta}). \quad (4.3)$$

Although the fits are performed with an unbinned likelihood, the fit quality is also quantified by a reduced χ^2 , computed by comparing the fitted model with the binned distribution shown in the corresponding figure, with the number of degrees of freedom (ndof) given by the number of bins minus the number of free parameters.

The two shapes are then combined into a signal-plus-background model, whose extended likelihood and parametrisation in terms of the branching fraction are described in Section 4.4.

The statistical behaviour of the procedure is investigated using pseudo-experiments generated with toy-MC. These studies include the evaluation of the pull distributions, providing checks of the agreement between the fitted and injected signal

contributions and of the reliability of the estimated uncertainties.

4.2 Signal model

The signal is modelled with a Double-sided Crystal Ball (DCB) probability density function. This choice is motivated by the shape of the reconstructed $m_{e\mu}$ distribution in the simulated $Z \rightarrow e\mu$ sample. The $Z \rightarrow e\mu$ decay produces a narrow peak around the Z boson mass, whose width results from the convolution of the natural width of the Z boson ($\Gamma_Z \simeq 2.5 \text{ GeV}$) with the experimental resolution on the lepton momenta. Radiative losses, namely final-state radiation and, for electrons, bremsstrahlung in the detector material, are not fully accounted for in the reconstructed momenta. Since these effects can only reduce the measured energy, they shift the peak slightly below m_Z and produce a low-mass tail, so that the resulting distribution is asymmetric. The Crystal Ball function combines a Gaussian core with an independent power-law tail on each side of the peak, and is therefore suited to describe these effects.

The invariant mass distribution of the simulated signal after the full event selection, together with the corresponding DCB fit, is shown in Figure 4.1. The fit describes the narrow peak and the non-Gaussian tails of the distribution, with a $\chi^2/\text{ndof} = 1.58$.

The DCB is characterised by six parameters: the peak position μ , the width σ of the Gaussian core, and two pairs of parameters, (α_L, n_L) and (α_R, n_R) , controlling the transition from the Gaussian core to the left and right power-law tails, respectively [65]. Its expression is

$$f_{\text{sig}}(m_{e\mu}; \mu, \sigma, \alpha_L, n_L, \alpha_R, n_R) = \mathcal{N} \cdot \begin{cases} A_L (B_L - t)^{-n_L}, & \text{if } t < -\alpha_L \\ \exp\left(-\frac{t^2}{2}\right), & \text{if } -\alpha_L \leq t \leq \alpha_R \\ A_R (B_R + t)^{-n_R}, & \text{if } t > \alpha_R \end{cases} \quad (4.4)$$

where $t = (m_{e\mu} - \mu)/\sigma$ and $\mathcal{N}$ is a normalisation constant, such that the integral of f_{sig} over the fit range is equal to one. The constants $A_{L,R}$ and $B_{L,R}$ are not free parameters: they are fixed by requiring the continuity of the function and of its first derivative at $t = -\alpha_L$ and $t = \alpha_R$, which gives




FIGURE 4.1: Invariant mass distribution of the simulated $Z \rightarrow e\mu$ signal after the full event selection, with the result of the DCB fit superimposed. The fitted values of the DCB parameters and the reduced χ^2 are reported in the figure.

$$A_{L,R} = \left(\frac{n_{L,R}}{|\alpha_{L,R}|} \right)^{n_{L,R}} \exp \left(-\frac{|\alpha_{L,R}|^2}{2} \right), \quad B_{L,R} = \frac{n_{L,R}}{|\alpha_{L,R}|} - |\alpha_{L,R}|. \quad (4.5)$$

The parameters are determined by fitting the simulated signal sample after the full event selection, over the range $70 \text{ GeV} < m_{e\mu} < 110 \text{ GeV}$. The fitted values of the parameters are reported in Fig. 4.1. The peak position, $\mu = 90.76 \pm 0.06 \text{ GeV}$, lies slightly below m_Z , and the left tail starts closer to the core than the right one ($\alpha_L < \alpha_R$), as expected from the radiative losses discussed above.

Since the number of signal events in data is expected to be small, I rely on this simulation fit to fully characterize the signal line-shape. Therefore, in the following signal-plus-background fit, these parameters are kept fixed, so that the signal shape enters the final model without additional degrees of freedom.

4.3 Background model

The background distribution is modelled using the simulated samples described in Section 3.1.2. After the complete event selection, the contributions from the




FIGURE 4.2: Invariant mass distribution of the simulated total background samples after the full event selection, with the result of the third-order Chebyshev polynomial fit superimposed. The fitted values of the Chebyshev parameters and the reduced χ^2 are reported in the figure.

$Z \rightarrow ee$, $Z \rightarrow \mu\mu$, $Z \rightarrow \tau\tau$, WW , $t\bar{t}$, and WZ processes are combined into a single weighted dataset. This provides the background distribution in the $m_{e\mu}$ variable used to determine the analytical model.

The resulting distribution is shown in Fig. 4.2. In the mass region considered, the distribution is smooth and does not exhibit the narrow resonance characteristic of the signal. This motivates the use of a polynomial function to describe the overall background shape and a Chebyshev polynomial was chosen [66]. Polynomials of different order were tested, and the third-order polynomial was selected as the nominal model, as it yields the best reduced χ^2 among the tested functions.

The background PDF is described by:

$$f_{\text{bkg}}(m_{e\mu}) = \frac{1}{\mathcal{N}} [1 + c_1 T_1(x) + c_2 T_2(x) + c_3 T_3(x)], \quad (4.6)$$

where $\mathcal{N}$ is the normalisation factor and c_1 , c_2 , and c_3 are the free coefficients of the polynomial. The variable x is obtained by mapping, according to [67], the invariant-mass range onto the interval $[-1, 1]$,

$$x = \frac{2m_{e\mu} - (m_{\max} + m_{\min})}{m_{\max} - m_{\min}}$$

where $m_{\min} = 73 \text{ GeV}$ and $m_{\max} = 110 \text{ GeV}$. The first three Chebyshev polynomials are

$$\begin{aligned} T_1(x) &= x, \\ T_2(x) &= 2x^2 - 1, \\ T_3(x) &= 4x^3 - 3x. \end{aligned}$$

The background fit is performed in the range $73 \text{ GeV} < m_{e\mu} < 110 \text{ GeV}$. The lower boundary excludes the low-mass plateau visible in the simulated $Z \rightarrow \tau\tau$ contribution in the invariant-mass stack plot shown in Fig. 3.5. Restricting the fit range in this way allows the polynomial to describe the smooth background component in the region relevant to the $Z \rightarrow e\mu$ signal without requiring a higher polynomial order.

The analytical model is fitted to the combined weighted dataset with an unbinned maximum-likelihood fit. The fitted coefficients and the corresponding fit result are shown in Figure 4.2. The reduced χ^2 calculated from the binned distribution is used as an additional measure of the agreement between the model and the simulated background.

The resulting background PDF is then combined with the signal model to build the complete signal-plus-background model. In this fit, the background coefficients are left free, and the simulation provides only their starting values.

4.4 Total fit and validation

The signal and background shapes determined in the previous sections are now combined into a single model, which is validated on simulation before being applied to data. The model is first fitted to a simulated dataset with a known signal contribution, and the signal yield is expressed in terms of the branching fraction. Pseudo-experiments are then used to verify that the estimator is unbiased and that its uncertainty is correctly evaluated, both in the presence and in the absence of signal.

4.4.1 Combined model and branching fraction

The signal and background models are combined into a single model,

$$f(m_{e\mu}; \boldsymbol{\theta}) = \frac{N_{\text{sig}} f_{\text{sig}}(m_{e\mu}) + N_{\text{bkg}} f_{\text{bkg}}(m_{e\mu})}{N_{\text{sig}} + N_{\text{bkg}}}, \quad (4.7)$$

where N_{sig} and N_{bkg} are the signal and background yields.

Since the total number of events is also determined by the fit, the likelihood is extended with a Poisson term,

$$\mathcal{L}_{\text{ext}}(\boldsymbol{\theta}) = \frac{\nu^N e^{-\nu}}{N!} \prod_{i=1}^N \frac{[N_{\text{sig}} f_{\text{sig}}(m_{e\mu,i}) + N_{\text{bkg}} f_{\text{bkg}}(m_{e\mu,i})]}{N_{\text{sig}} + N_{\text{bkg}}}, \quad (4.8)$$

where N is the observed number of events and $\nu = N_{\text{sig}} + N_{\text{bkg}}$ is the number expected from the model. The fit therefore constrains both the shapes and the yields of the two components.

The model is first tested on a simulated dataset, obtained by summing the weighted background samples and the signal sample, the latter normalised to a branching fraction of 10^{-6} . The fit is performed in the range $73 < m_{e\mu} < 110$ GeV. The parameters describing the signal shape are fixed to the values determined in Section 4.2, so that the only free parameter of the signal component is its yield N_{sig} . The background yield N_{bkg} and the coefficients of the Chebyshev polynomial are free.

The result is shown in Figure 4.3. The fitted signal yield is $N_{\text{sig}} = 877 \pm 53$, with an expected 904, and the fitted background yield is $N_{\text{bkg}} = 14324 \pm 98$, with 14297 expected from the simulation in the fit range. Both yields agree with the expectation within one standard deviation, -0.5σ and $+0.3\sigma$ respectively, and the reduced χ^2 is 1.55.

The signal yield is then expressed in terms of the branching fraction:

$$N_{\text{sig}} = \mathcal{B}(Z \rightarrow e\mu) \sigma_Z L_{\text{int}} \varepsilon, \quad (4.9)$$

where σ_Z is the inclusive Z boson production cross section, $L_{\text{int}} = 109.95 \text{ fb}^{-1}$ is the integrated luminosity and ε is the signal efficiency. The efficiency ε corresponds to the total selection efficiency defined in Eq. (3.5), with two differences: it is restricted to the fit range, and each simulated event is counted with its weight, which includes the corrections to the simulated lepton efficiencies and the pile-up




FIGURE 4.3: Invariant mass distribution of the simulated signal-plus-background, with the result of the composed model fit superimposed.

reweighting of Eq. (3.1). It is therefore evaluated as the ratio between the weighted number of signal events passing the full selection in the fit range and the number of $Z \rightarrow e\mu$ decays expected for the injected branching fraction,

$$\varepsilon = \frac{N_{\text{sig}}}{\sigma_Z L_{\text{int}} B_{\text{inj}}} = 0.1369.$$

The three quantities σ_Z , L_{int} and ε are constant in the fit, so that $\mathcal{B}(Z \rightarrow e\mu)$ replaces N_{sig} as the free parameter of the signal component. Their uncertainties are not included at this stage.

The branching fraction is left free in the range $[-10^{-4}, +10^{-4}]$. Negative values, although unphysical, are allowed so that the estimator is not constrained by a boundary at zero. With this parametrisation the fit to the same dataset converges and gives

$$\mathcal{B} = (9.70 \pm 0.59) \times 10^{-7}, \quad (4.10)$$

which is 0.5 standard deviations below the injected value. The fitted curve and the reduced χ^2 are those of the previous fit, since the two parametrisations are equivalent.

The fit is repeated on simulated datasets in which the signal is normalised to $\mathcal{B}_{\text{inj}} = 10^{-7}$ and $\mathcal{B}_{\text{inj}} = 0$. The fitted values, reported in Table 4.1, agree with the injected ones within one standard deviation. The properties of the estimator are studied with pseudo-experiments in the next section.

$\mathcal{B}_{\text{inj}}$	$\mathcal{B}_{\text{fit}}$	pull mean	pull width
1.00×10^{-6}	$(9.70 \pm 0.59) \times 10^{-7}$	0.006 ± 0.031	0.992 ± 0.022
1.00×10^{-7}	$(0.91 \pm 0.50) \times 10^{-7}$	-0.023 ± 0.031	0.974 ± 0.022
0	$(-0.06 \pm 0.50) \times 10^{-7}$	-0.020 ± 0.031	0.989 ± 0.022

TABLE 4.1: Fitted branching fraction obtained from the fit to the simulated dataset for each injected value, and mean and width of the Gaussian fitted to the corresponding pull distribution, obtained from 1000 pseudo-experiments.

4.4.2 Validation with pseudo-experiments

Pseudo-experiments are generated for the three values of $\mathcal{B}_{\text{inj}}$ considered above. The last two probe the regime of a small or absent signal, in which the search is performed.

For each value, 1000 pseudo-datasets are generated from the signal-plus-background model, with the signal yield set by $\mathcal{B}_{\text{inj}}$ and the background yield and shape set to the values obtained from the nominal fit. In each pseudo-experiment the number of events is drawn from a Poisson distribution with mean equal to the expected total yield, and the $m_{e\mu}$ values are sampled from the model PDF. Each pseudo-dataset is then fitted with the same model used for the nominal fit: the signal shape, σ_Z , L_{int} and ε are kept fixed, whereas $\mathcal{B}$, the background yield and all the background shape parameters are left free.


For every pseudo-experiment the pull is computed,

$$\frac{\mathcal{B}_{\text{fit}} - \mathcal{B}_{\text{inj}}}{\sigma_{\mathcal{B}_{\text{fit}}}}, \quad (4.11)$$


where $\sigma_{\mathcal{B}_{\text{fit}}}$ is the uncertainty on $\mathcal{B}_{\text{fit}}$ returned by the fit. If the estimator were unbiased and its uncertainty were correctly evaluated, the pull would follow a Gaussian distribution with zero mean and unitary standard deviation.

Figures 4.4, 4.5 and 4.6 show, for each value of $\mathcal{B}_{\text{inj}}$, the distributions for each pseudo-experiment of the fitted branching fraction, of its uncertainty, of the pull


and of the minimum of $-\ln \mathcal{L}$. The pull distributions are fitted with a Gaussian function, whose mean and width are reported in Table 4.1. In all cases the fitted branching fraction is centred on the injected value, and the distribution of the minimum of $-\ln \mathcal{L}$ exhibits a regular, approximately Gaussian shape, as expected if all fits converge. The mean of the pull is compatible with zero within one standard deviation in all cases, and its width deviates from one by at most 1.2 standard deviations, for $\mathcal{B}_{\text{inj}} = 10^{-7}$. The estimator is therefore unbiased and its uncertainty is correctly evaluated, and the model can be applied to data.



(a) Fitted branching fraction.



(b) Uncertainty on the fitted branching fraction.



(c) Pull, with a Gaussian fit superimposed.


(d) Minimum of $-\ln \mathcal{L}$.

FIGURE 4.4: Distributions obtained from 1000 pseudo-experiments for an injected branching fraction $\mathcal{B}_{\text{inj}} = 10^{-6}$.

4.5 Fit to data

The studies of Section 4.4 show that the fit returns an unbiased estimate of $\mathcal{B}(Z \rightarrow e\mu)$ with a reliable uncertainty. The region $80 < m_{e\mu} < 100$ GeV is




FIGURE 4.5: Distributions obtained from 1000 pseudo-experiments for an injected branching fraction $\mathcal{B}_{\text{inj}} = 10^{-7}$.

therefore unblinded. Figure 4.7 shows the $m_{e\mu}$ distribution after unblinding, compared with the simulated background. The signal expected for $\mathcal{B}(Z \rightarrow e\mu) = 10^{-6}$ is superimposed for illustration. The data follow the smooth background shape over the whole range, with no visible enhancement at the Z boson mass. This comparison is qualitative: in the fit, the background shape and yield are determined from data.

The fit is applied to the 2024 data sample in the range $73 < m_{e\mu} < 110$ GeV. The model is that of Eq. (4.7), in the configuration used for the validation. The signal shape is fixed to the values obtained from simulation (Section 4.2), and σ_Z , L_{int} and ε are constant, so that $\mathcal{B}$ is the only free parameter of the signal component. The background yield and the three Chebyshev coefficients are free and are determined directly from data. The simulation provides only their starting values. Data events are unweighted, hence the extended likelihood of Eq. (4.8) is




FIGURE 4.6: Distributions obtained from 1000 pseudo-experiments generated under the background-only hypothesis, $\mathcal{B}_{\text{inj}} = 0$.

used without any weight correction.

The result is shown in Figure 4.8. The fit converges and the reduced χ^2 is 0.98. The fitted background yield is $N_{\text{bkg}} = 15010 \pm 153$, and the branching fraction is

$$\mathcal{B}(Z \rightarrow e\mu) = (-1.06 \pm 1.00) \times 10^{-7},$$

where the uncertainty is statistical only.

The negative value, compatible with zero within about one standard deviation, indicates no evidence for a $Z \rightarrow e\mu$ signal. The statistical interpretation of this result is presented in Chapter 5.




FIGURE 4.7: Invariant mass distribution of the selected $e\mu$ pairs in data (points, with statistical uncertainties) after unblinding, compared with the simulated backgrounds (stacked histograms). The hatched histogram shows the signal expected for $\mathcal{B}(Z \rightarrow e\mu) = 10^{-6}$. The lower panel shows the ratio between data and the total simulated background.




FIGURE 4.8: Invariant mass distribution of the collision data after the full event selection, with the result of the signal-plus-background fit superimposed. The fitted parameters and the reduced χ^2 are reported in the figure.

Chapter 5

Statistical interpretation and results

The fit to data presented in Section 4.5 shows no evidence of a $Z \rightarrow e\mu$ signal: the fitted branching fraction is compatible with zero within about one standard deviation. This result is quantified in two steps. Section 5.1 evaluates the compatibility of the data with the background-only hypothesis, and Section 5.2 derives an upper limit on $\mathcal{B}(Z \rightarrow e\mu)$ at 95% confidence level (CL).

The statistical interpretation follows the frequentist procedure recommended by the LHC Higgs Combination Group [68], with test statistics based on the profile likelihood ratio [69]. The parameter of interest, i.e., the quantity on which the inference is performed, is the branching fraction $\mathcal{B}(Z \rightarrow e\mu)$, which determines the signal yield through Eq. (4.9). The background yield and the Chebyshev coefficients are nuisance parameters, denoted collectively by $\boldsymbol{\theta}$: they are not of direct interest, but must be determined from the data. As in Chapter 4, σ_Z , L_{int} , and ε are kept fixed, so that the likelihood coincides with the extended likelihood of Eq. (4.8) and the results include statistical uncertainties only.

For a tested value of $\mathcal{B}$, the profile likelihood ratio is

$$\lambda(\mathcal{B}) = \frac{\mathcal{L}(\mathcal{B}, \hat{\boldsymbol{\theta}}_{\mathcal{B}})}{\mathcal{L}(\hat{\mathcal{B}}, \hat{\boldsymbol{\theta}})}, \quad (5.1)$$

where $\hat{\mathcal{B}}$ and $\hat{\boldsymbol{\theta}}$ maximise the likelihood, and $\hat{\boldsymbol{\theta}}_{\mathcal{B}}$ are the values of the nuisance parameters that maximise it for a given value of $\mathcal{B}$. By construction $0 \leq \lambda \leq 1$, and values close to one indicate that the tested $\mathcal{B}$ describes the data almost as well as the best fit ($\hat{\mathcal{B}}$). As in the fit of Section 4.5, $\hat{\mathcal{B}}$ is not restricted to positive values.

5.1 Background-only hypothesis test

A signal would appear as an excess of events at the Z boson mass. The test quantifies the probability that the background alone produced data at least as incompatible with the background-only hypothesis as the observed ones. The test statistic is

$$t_0 = -2 \ln \lambda(0). \quad (5.2)$$

For a large sample, Wilks' theorem [69] gives the significance and the corresponding p-value as

$$Z = \sqrt{t_0}, \quad p_0 = 1 - \Phi(Z), \quad (5.3)$$

where Φ is the cumulative distribution of the standard normal distribution. Evidence and observation of a signal conventionally require $Z \geq 3$ and $Z \geq 5$, respectively.

The fit of Section 4.5 gives $\hat{\mathcal{B}} = (-1.06 \pm 1.00) \times 10^{-7}$, and Eq. (5.2) gives $t_0 = 1.13$, corresponding to $p_0 = 0.14$ and $Z = 1.06$: the data show no excess over the background-only expectation.

5.2 Upper limit

Since the data show no excess, I can set an upper limit on the branching fraction. The test of Section 5.1 evaluates the compatibility of the data with the background-only hypothesis. Here the roles of the two hypotheses are exchanged: a signal with a given branching fraction is assumed to be present, and the probability is evaluated that a downward fluctuation produces data at least as incompatible with that signal as the observed ones. If this probability is small, the assumed branching fraction is excluded. The procedure is repeated for a set of values of $\mathcal{B}$, and the upper limit is the value above which all branching fractions are excluded at 95% CL.

For a tested value of $\mathcal{B}$, the test statistic is

$$q_{\mathcal{B}} = \begin{cases} -2 \ln \lambda(\mathcal{B}) & \hat{\mathcal{B}} \leq \mathcal{B}, \\ 0 & \hat{\mathcal{B}} > \mathcal{B}, \end{cases} \quad (5.4)$$

with $\lambda(\mathcal{B})$ defined in Eq. (5.1). Data preferring a larger branching fraction than the tested one are not regarded as evidence against it, so that the resulting interval is one-sided. Larger values of $q_{\mathcal{B}}$ correspond to increasing incompatibility between the data and the tested $\mathcal{B}$.

Unlike the test of Section 5.1, the distributions of $q_{\mathcal{B}}$ are not taken from asymptotic formulae, but are obtained with pseudo-experiments, which do not rely on the large-sample approximation. For each tested $\mathcal{B}$, 10^4 pseudo-datasets are generated under the signal-plus-background hypothesis with that branching fraction, and 10^4 under the background-only hypothesis. In both cases the nuisance parameters are set to the values that maximise the likelihood of the observed data for the corresponding hypothesis. Each pseudo-dataset is fitted, and $q_{\mathcal{B}}$ is computed as for the data. This gives the distributions $g(q_{\mathcal{B}} | \mathcal{B})$ and $g(q_{\mathcal{B}} | 0)$, from which

$$\text{CL}_{s+b} = \int_{q_{\mathcal{B},\text{obs}}}^{\infty} g(q_{\mathcal{B}} | \mathcal{B}) dq_{\mathcal{B}}, \quad \text{CL}_b = \int_{q_{\mathcal{B},\text{obs}}}^{\infty} g(q_{\mathcal{B}} | 0) dq_{\mathcal{B}}, \quad (5.5)$$

where g denotes the probability density of the test statistic. CL_{s+b} is the probability of obtaining data at least as incompatible with the tested signal as those observed, if the signal is present. CL_b is the same probability under the background-only hypothesis.

To avoid excluding signals to which the analysis is not sensitive, as may happen with CL_{s+b} alone in the presence of a downward fluctuation of the background, the CL_s method [70, 71] is used:

$$\text{CL}_s = \frac{\text{CL}_{s+b}}{\text{CL}_b}, \quad (5.6)$$

and a value of $\mathcal{B}$ is excluded at 95% CL if $\text{CL}_s < 0.05$.

CL_s is evaluated at 15 values of $\mathcal{B}$ between 0 and 4×10^{-7} . The upper limit is the value at which $\text{CL}_s = 0.05$, obtained by interpolation between adjacent points. The expected limit is obtained by replacing $q_{\mathcal{B},\text{obs}}$ with the median of $g(q_{\mathcal{B}} | 0)$, i.e. with the outcome of a typical background-only experiment. Replacing it with the quantiles of $g(q_{\mathcal{B}} | 0)$ at 15.87% and 84.13% (2.28% and 97.72%), i.e. the fractions corresponding to $\pm 1\sigma$ ($\pm 2\sigma$) of a Gaussian distribution, gives the bands containing 68% (95%) of the limits expected in the absence of a signal.

The scan is shown in Fig. 5.1 and the limits are summarised in Table 5.1. The observed upper limit is

$$\mathcal{B}(Z \rightarrow e\mu) < 1.4 \times 10^{-7} \quad \text{at 95\% CL}, \quad (5.7)$$

to be compared with an expected limit of 2.0×10^{-7} .




FIGURE 5.1: CL_s (black points) and CL_{s+b} (blue points) as a function of the tested branching fraction $\mathcal{B}(Z \rightarrow e\mu)$. The dashed line shows the median expected CL_s under the background-only hypothesis, and the green and yellow bands its 68% and 95% ranges. The red line marks $\text{CL}_s = 0.05$.

	$\mathcal{B}(Z \rightarrow e\mu) [10^{-7}]$
Observed	1.4
Expected (median)	2.0
Expected $\pm 1\sigma$	[1.4, 2.8]
Expected $\pm 2\sigma$	[1.1, 3.7]

TABLE 5.1: Observed and expected upper limits at 95% CL on $\mathcal{B}(Z \rightarrow e\mu)$, statistical uncertainties only.

5.3 Comparison with the Run 2 result

The CMS Run 2 analysis reports an observed (expected) upper limit of $1.9 (2.0) \times 10^{-7}$, obtained with an integrated luminosity of 138 fb^{-1} at 13 TeV [1]. The comparison is performed between expected limits, which quantify the sensitivity of the analysis independently of the fluctuations in the observed data. The two are almost identical, although the present analysis uses about 20% less integrated luminosity and a cut-based selection, whereas the Run 2 analysis classified events with a boosted decision tree. The comparison is nevertheless only indicative: the Run 2 limit includes the full set of systematic uncertainties, while the present one is statistical only, and the slightly larger Z boson production cross section at 13.6 TeV partially compensates the smaller dataset.

In the Run 2 analysis, the largest contribution to the systematic uncertainty on the branching fraction, 7.9×10^{-8} , arose from the parameters of the background functions, determined from the data, while the statistical component amounted to 5.6×10^{-8} [1]. Since these parameters are also left free in the present fit, the uncertainty obtained here, 1.00×10^{-7} , is to be compared with the quadrature sum of the two terms, 0.97×10^{-7} , to which it is closely comparable. The remaining systematic contributions were subdominant in the Run 2 analysis, but must nevertheless be included before the result can be regarded as final.

Conclusions

The Standard Model (SM) has been confirmed by every direct test performed so far at colliders, yet the observation of neutrino oscillations shows that lepton flavour is not an exact symmetry of nature. In the charged sector, however, flavour violation induced by neutrino masses remains far beyond experimental reach: for the decay $Z \rightarrow e\mu$, the predicted branching fraction is of order 10^{-60} . Several extensions of the SM instead predict rates that are potentially accessible at the LHC. Any observation of charged lepton flavour violation (CLFV) would therefore constitute unambiguous evidence of new physics.

This thesis has presented a search for $Z \rightarrow e\mu$ with proton–proton collisions recorded by the CMS experiment in 2024, during Run 3 of the LHC (2022–2026). The data were collected at $\sqrt{s} = 13.6$ TeV, corresponding to an integrated luminosity of 109.95 fb^{-1} . The analysis was first developed and validated on simulation. The event selection isolates an opposite-charge electron–muon pair compatible with the decay of a Z boson, and suppresses the Drell–Yan, diboson and top-quark pair backgrounds through requirements on lepton quality, b -tagged jets and missing transverse momentum. The signal is searched for as a narrow peak at the Z boson mass over a smooth background in the $e\mu$ invariant mass spectrum, through a fit that measures the branching fraction directly. Pseudo-experiments showed that the fit returns an unbiased estimate with a correctly evaluated uncertainty, both in the presence and in the absence of signal.

Applied to data, the fit yields a best estimate of the branching fraction of

$$\mathcal{B}(Z \rightarrow e\mu) = (-1.06 \pm 1.00) \times 10^{-7},$$

a value compatible with zero within about one standard deviation: no evidence for a $Z \rightarrow e\mu$ signal is observed. An upper limit on this parameter is therefore set with the CL_s method, evaluated with pseudo-experiments,

$$\mathcal{B}(Z \rightarrow e\mu) < 1.4 \times 10^{-7} \quad \text{at 95\% CL},$$

against a median expected limit of 2.0×10^{-7} , with a 68% band of $[1.4, 2.8] \times 10^{-7}$. The comparison with the CMS Run 2 (2015–2018) analysis [1], based on 138 fb^{-1} at 13 TeV, is performed between expected limits. The Run 2 expected limit, 2.0×10^{-7} , is identical to the one obtained here, although the present analysis relies on about 20% less integrated luminosity and on a cut-based selection instead of a boosted decision tree. Since the present result includes statistical uncertainties only, the remaining systematic contributions, subdominant in Run 2, must still be included before it can be regarded as final. For CMS, the main implication is that the 2024 data alone already match the sensitivity of the complete Run 2 dataset, placing the Run 3 analysis in a favourable position to set the most stringent direct limit on this decay.

Once the systematic uncertainties are included, the analysis can be extended to the full Run 3 dataset. As long as the search remains statistically dominated, the expected limit scales approximately as $1/\sqrt{L_{\text{int}}}$, where L_{int} is the integrated luminosity: with about 350 fb^{-1} , it would reach about 1.1×10^{-7} , well below the current best limit. Further gains can come from the analysis strategy, through a multivariate classification of the events and a dedicated study of the background parametrisation. The same framework can then be projected to the High-Luminosity LHC, where about 3000 fb^{-1} are foreseen, and to future facilities such as the Future Circular Collider and a Muon Collider.

The search will also be extended to the remaining channels, $Z \rightarrow e\tau$ and $Z \rightarrow \mu\tau$, whose most stringent limits, 5.0×10^{-6} and 6.5×10^{-6} , were set by the ATLAS experiment [18]. These channels are experimentally more challenging: the neutrinos produced in the τ decays prevent the full reconstruction of the Z boson, so that the signal no longer appears as a narrow peak. As in the CMS Run 2 analysis, the signal will be separated from the background with a multivariate approach. Systematic uncertainties will play a larger role than in the $e\mu$ channel; precisely because these channels are less constrained, they offer considerable room for improvement. Together with $Z \rightarrow e\mu$, they will complete the search for CLFV in all decays of the Z boson into two charged leptons of different flavour. Each improvement of these limits further restricts the room left for new physics in the lepton sector, or brings closer the moment in which it is revealed.

Bibliography

- [1] CMS Collaboration. “Search for flavor violating decays of the Z boson to charged leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* (2025). DOI: 10.1103/81w4-6lyk.
- [2] M. Thomson. *Modern Particle Physics*. Cambridge University Press, 2013.
- [3] D. J. Griffiths. *Introduction to elementary particles*. Physics textbook. New York, NY: Wiley, 2008.
- [4] D. Rebuzzi. *Dispense del corso di Introduzione alla Fisica Nucleare e Sub-nucleare*. Materiale didattico non pubblicato. 2025.
- [5] G. Arnison et al. “Experimental observation of isolated large transverse energy electrons with associated missing energy at $s=540$ GeV”. In: *Physics Letters B* 122.1 (1983), pp. 103–116. DOI: 10.1016/0370-2693(83)91177-2.
- [6] M. Banner et al. “Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN $p\bar{p}$ collider”. In: *Physics Letters B* 122.5 (1983), pp. 476–485. DOI: 10.1016/0370-2693(83)91605-2.
- [7] G. Arnison et al. “Experimental observation of lepton pairs of invariant mass around 95 GeV/ c^2 at the CERN SPS collider”. In: *Physics Letters B* 126.5 (1983), pp. 398–410. DOI: 10.1016/0370-2693(83)90188-0.
- [8] P. Bagnaia et al. “Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $p\bar{p}$ collider”. In: *Physics Letters B* 129.1 (1983), pp. 130–140. DOI: 10.1016/0370-2693(83)90744-X.
- [9] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321.

-
- [10] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508.
 - [11] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. In: *Phys. Rev. Lett.* 13 (20 Nov. 1964), pp. 585–587. DOI: 10.1103/PhysRevLett.13.585.
 - [12] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716.1 (2012).
 - [13] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1 (2012).
 - [14] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Physical Review Letters* 81.8 (1998).
 - [15] Q. R. Ahmad et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Physical Review Letters* 89.1 (2002).
 - [16] J. I. Illana, M. Jack, and T. Riemann. “Predictions for $Z \rightarrow \mu\tau$ and related reactions”. In: (2000).
 - [17] Y. Kuno and Y. Okada. “Muon decay and physics beyond the standard model”. In: *Reviews of Modern Physics* 73.1 (Jan. 2001), pp. 151–202. DOI: 10.1103/revmodphys.73.151.
 - [18] ATLAS Collaboration. “Search for Lepton-Flavour Violation in Z -Boson Decays with τ Leptons with the ATLAS Detector”. In: *Phys. Rev. Lett.* 127 (2021).
 - [19] Particle Data Group. “Review of Particle Physics”. In: *Int. J. Mod. Phys. A* 41 (2026).
 - [20] V. Chekhovsky et al. “Measurements of the inclusive W and Z boson production cross sections and their ratios in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”. In: *Journal of High Energy Physics* 2026.1 (Jan. 2026). DOI: 10.1007/jhep01(2026)047.
 - [21] CMS Collaboration. *The End of Run 3*. 2026. URL: <https://cms.cern/news/end-run-3> (visited on 07/09/2026).

- [22] O. S. Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-1.
- [23] Forthommel. *Overview of current particle and nuclear facilities at the CERN accelerator complex*. Wikimedia Commons, license CC BY-SA 3.0. 2012. URL: <https://commons.wikimedia.org/wiki/File:Cern-accelerator-complex.svg> (visited on 07/14/2026).
- [24] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010.
- [25] CERN. *HL-LHC Project Schedule*. 2026. URL: <https://project-hl-lhc-industry.web.cern.ch/content/project-schedule> (visited on 08/05/2026).
- [26] ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.
- [27] CMS Collaboration. “The CMS Experiment at the CERN LHC”. In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.
- [28] LHCb Collaboration. “The LHCb Detector at the LHC”. In: *JINST* 3 (2008), S08005. DOI: 10.1088/1748-0221/3/08/S08005.
- [29] ALICE Collaboration. “The ALICE Experiment at the CERN LHC”. In: *JINST* 3 (2008), S08002. DOI: 10.1088/1748-0221/3/08/S08002.
- [30] CMS Collaboration. *Detector*. CMS Experiment, CERN. URL: <https://cms.cern/detector> (visited on 07/15/2026).
- [31] I. Neutelings. *3D Coordinate System for CMS*. URL: https://tikz.net/axis3d_cms/ (visited on 08/05/2026).
- [32] CMS Collaboration. “Description and performance of track and primary-vertex reconstruction with the CMS tracker”. In: *JINST* 9 (2014), P10009. DOI: 10.1088/1748-0221/9/10/P10009.
- [33] CMS Collaboration. *The Phase-2 Upgrade of the CMS Tracker*. Tech. rep. CERN-LHCC-2017-009, CMS-TDR-014. CERN, 2017.
- [34] CMS Collaboration. *The CMS electromagnetic calorimeter project: Technical Design Report*. Tech. rep. CERN-LHCC-97-033, CMS-TDR-4. CERN, 1997.

- [35] CMS Collaboration. “Energy calibration and resolution of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 7$ TeV”. In: *JINST* 8 (2013), P09009. DOI: 10.1088/1748-0221/8/09/P09009.
- [36] CMS Collaboration. *The CMS hadron calorimeter project: Technical Design Report*. Tech. rep. CERN-LHCC-97-031, CMS-TDR-2. CERN, 1997.
- [37] CMS Collaboration. *The Phase-2 Upgrade of the CMS Endcap Calorimeter*. Tech. rep. CERN-LHCC-2017-023, CMS-TDR-019. CERN, 2017.
- [38] CMS Collaboration. *The CMS muon project: Technical Design Report*. Tech. rep. CERN-LHCC-97-032, CMS-TDR-3. CERN, 1997.
- [39] A. Colaleo et al. *CMS Technical Design Report for the Muon Endcap GEM Upgrade*. Tech. rep. 2015.
- [40] CMS Collaboration. *The Phase-2 Upgrade of the CMS Muon Detectors*. Tech. rep. CERN-LHCC-2017-012, CMS-TDR-016. CERN, 2017.
- [41] CMS Collaboration. *CMS TriDAS project: Technical Design Report, Volume 1: The Trigger Systems*. Tech. rep. CERN-LHCC-2000-038, CMS-TDR-6-1. CERN, 2000.
- [42] CMS Collaboration. “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JINST* 15 (2020), P10017. DOI: 10.1088/1748-0221/15/10/P10017.
- [43] CMS Collaboration. “The CMS high level trigger”. In: *Eur. Phys. J. C* 46.3 (2006), pp. 605–667. DOI: 10.1140/epjc/s2006-02495-8.
- [44] CMS Collaboration. “Particle-flow reconstruction and global event description with the CMS detector”. In: *JINST* 12.10 (2017), P10003. DOI: 10.1088/1748-0221/12/10/P10003.
- [45] CMS Collaboration. “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”. In: *JINST* 16 (2021), P05014. DOI: 10.1088/1748-0221/16/05/P05014.
- [46] M. Seidel. “Physics Objects in CMS Run 3”. In: (2026). for the CMS Collaboration; talk at the 18th International Workshop on Top Quark Physics (Top2025).

-
- [47] CMS Collaboration. *Finding electrons and photons with the CMS detector*. Feb. 3, 2021. URL: <https://cms.cern/news/finding-electrons-and-photons-cms-detector> (visited on 08/11/2026).
- [48] CMS Collaboration. *Electron and photon reconstruction and identification performance at CMS in 2022 and 2023*. Tech. rep. CMS-DP-2024-052. CERN, July 2024.
- [49] CMS Open Data. *Jets*. URL: <https://cms-opendata-guide.web.cern.ch/analysis/selection/objects/jets/> (visited on 08/11/2026).
- [50] M. Cacciari, G. P. Salam, and G. Soyez. “The anti-kt jet clustering algorithm”. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063.
- [51] CMS Collaboration. “Run 3 commissioning results of heavy-flavor jet tagging at $\sqrt{s} = 13.6$ TeV with CMS data using a modern framework for data processing”. In: (2024).
- [52] CMS Collaboration. “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector”. In: *Journal of Instrumentation* 14.07 (2019), P07004–P07004. DOI: 10.1088/1748-0221/14/07/p07004.
- [53] CMS Open Data. *Missing transverse momentum*. URL: <https://cms-opendata-guide.web.cern.ch/analysis/selection/objects/met/> (visited on 08/11/2026).
- [54] D. Bertolini et al. “Pileup per particle identification”. In: *Journal of High Energy Physics* 2014.10 (Oct. 2014). DOI: 10.1007/jhep10(2014)059.
- [55] CMS Collaboration. “Performance of the reconstruction and identification of high-momentum muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of Instrumentation* 15.02 (2020), P02027. DOI: 10.1088/1748-0221/15/02/P02027.
- [56] CMS Collaboration. “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of Instrumentation* 13.06 (2018), P06015. DOI: 10.1088/1748-0221/13/06/P06015.

- [57] CMS Collaboration. *Muon ID and Isolation efficiencies with muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV*. Tech. rep. CMS-DP-2024/067. CERN, July 2024.
- [58] C. Aimè. “Hidden sectors search at the CMS experiment and predictions for future colliders”. PhD thesis. Università degli studi di Pavia, Italy, Dec. 2022.
- [59] CMS Data Preservation & Open Access group. *CMS Open Data Workshop lessons*. 2026. URL: <https://cms-opendata-workshop.github.io/workshop2024-lesson-exploring-cms-nanoaod> (visited on 08/10/2026).
- [60] A. Tumasyan et al. “First measurement of the top quark pair production cross section in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”. In: *Journal of High Energy Physics* 2023.8 (Aug. 2023). DOI: 10.1007/jhep08(2023)204.
- [61] A. Hayrapetyan et al. “Measurement of inclusive and differential cross sections for W^+W^- production in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”. In: *Physics Letters B* 861 (2025), p. 139231. DOI: 10.1016/j.physletb.2024.139231.
- [62] M. Meena. *Measurement of WZ production cross section at 13.6 TeV*. Tech. rep. Geneva: CERN, 2025. DOI: 10.1007/978-981-95-1513-4_163.
- [63] CMS Collaboration. *A unified approach for jet tagging in Run 3 at $\sqrt{s} = 13.6$ TeV in CMS*. Tech. rep. CMS-DP-2024-066. CERN, 2024.
- [64] W. Verkerke and D. Kirkby. *The RooFit toolkit for data modeling*. 2003.
- [65] ROOT Collaboration. *RooCrystalBall Class Reference*. <https://root.cern.ch/doc/v636/classRooCrystalBall.html>. ROOT Data Analysis Framework Documentation. 2024.
- [66] ROOT Collaboration. *RooChebychev Class Reference, ROOT v6-36 Reference Guide*. <https://root.cern.ch/doc/v636/classRooChebychev.html>. ROOT Data Analysis Framework Documentation.
- [67] ROOT Collaboration. *RooFit MathFuncs.h Source File, ROOT v6-36 Reference Guide*. https://root.cern.ch/doc/v636/MathFuncs_8h_source.html. Function `chebychev`, transformation of the observable to the range $[-1, 1]$. Accessed: 2026-10-09.
- [68] *Procedure for the LHC Higgs boson search combination in Summer 2011*. Tech. rep. Geneva: CERN, 2011.

-
- [69] G. Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *The European Physical Journal C* 71.2 (Feb. 2011). DOI: 10.1140/epjc/s10052-011-1554-0.
- [70] T. Junk. “Confidence level computation for combining searches with small statistics”. In: *Nucl. Instrum. Meth. A* 434 (1999), pp. 435–443. DOI: 10.1016/S0168-9002(99)00498-2.
- [71] A. L. Read. “Presentation of search results: the CL_s technique”. In: *J. Phys. G* 28 (2002), pp. 2693–2704. DOI: 10.1088/0954-3899/28/10/313.