Search for Light Pseudoscalars in Exotic Decays of the Higgs Boson with the CMS Experiment

Dissertation

zur Erlangung des Doktorgrades

an der Fakultät für Mathematik,

Informatik und Naturwissenschaften

Fachbereich Physik

der Universität Hamburg

vorgelegt von

Lakshmi Priya Nair

aus

Kerala, Indien

Hamburg

2025

Declaration on oath

I hereby declare and affirm that this doctoral dissertation is my own work and that I have not used any aids and sources other than those indicated.

If electronic resources based on generative artificial intelligence (gAI) were used in the course of writing this dissertation, I confirm that my own work was the main and value-adding contribution and that complete documentation of all resources used is available in accordance with good scientific practice. I am responsible for any erroneous or distorted content, incorrect references, violations of data protection and copyright law, or plagiarism that may have been generated by the gAI.

Hamburg, den 10. Februar 2025 Lakshmi Priya Nair

Gutachter der Dissertation:	Dr. Alexei Raspereza Prof. Dr. Elisabetta Gallo
Zusammensetzung der Prüfungskommission:	Dr. Alexei Raspereza Prof. Dr. Elisabetta Gallo Dr. Katharina Behr Prof. Dr. Bernd A. Kniehl Prof. Dr. Johannes Haller
Vorsitzender der Prüfungskommission:	Prof. Dr. Bernd A. Kniehl
Datum der Disputation:	7. März 2025
Vorsitzender des Fach-Promotionsausschusses Physik:	Prof. Dr. Wolfgang J. Parak
Leiter des Fachbereichs Physik:	Prof. Dr. Markus Drescher
Dekan der Fakultät MIN:	Prof. Dr.-Ing. Norbert Ritter

Abstract

This thesis presents a search for a pair of light pseudoscalar bosons produced from the decays of the 125 GeV Higgs boson, focusing on final states with four tau leptons. This search is motivated by numerous beyond the standard model theories that predict an extended scalar sector, introducing a wide array of new physical states beyond the single Higgs boson observed in the standard model. Such theories, including the Two-Higgs-Doublet Model with an additional singlet (2HDM+S), suggest new interactions and couplings, enabling the 125 GeV Higgs to decay into lighter pseudoscalar bosons. These exotic decays serve as direct probes of new physics beyond the standard model. Given the enhanced couplings of light pseudoscalars to leptons in various scenarios within these models, the final states analyzed in this study are of particular relevance.

The search is based on data from proton-proton collisions at a center-of-mass energy of 13 TeV collected by the CMS detector at the LHC during the Run 2 data-taking period, corresponding to an integrated luminosity of $138\,\mathrm{fb}^{-1}$. The analysis targets pseudoscalars with masses between 4 and 15 GeV, which result in highly collimated decay products and present a challenging final state with overlapping fermion signatures. A specialized strategy, deviating from standard reconstruction and selection techniques, is implemented by utilizing muons and tracks for event identification.

No significant deviation of data with respect to the standard model expectation is observed in this search. Consequently, model-independent upper limits are set at 95% confidence level on the cross section times branching ratio for the 4τ final state, relative to the SM Higgs production cross section. Additionally, exclusion limits are established for various benchmark scenarios within the 2HDM+S framework, further constraining the parameter space of this model.

Zusammenfassung

Diese Dissertation präsentiert eine Suche nach einem Paar von leichten pseudoskalaren Bosonen, die aus den Zerfällen des 125 GeV Higgs-Bosons produziert werden, mit einem Fokus auf Endzustände mit vier Tau-Leptonen. Diese Suche wird durch zahlreiche Theorien jenseits des Standardmodells motiviert, die einen erweiterten skalarer Sektor vorhersagen und damit eine breite Palette neuer physikalischer Zustände über das im Standardmodell beobachtete einzelne Higgs-Boson hinaus einführen. Solche Theorien, einschließlich des Zwei-Higgs-Doublet-Modells mit einem zusätzlichen Singlet (2HDM+S), schlagen neue Interaktionen und Kopplungen vor, die es dem 125 GeV Higgs ermöglichen, in leichtere pseudoskalare Bosonen zu zerfallen. Diese exotischen Zerfälle dienen als direkte Sonden für neue Physik jenseits des Standardmodells. Angesichts der verstärkten Kopplungen von leichten Pseudoskalaren an Leptonen in verschiedenen Szenarien innerhalb dieser Modelle sind die in dieser Studie analysierten Endzustände von besonderer Relevanz.

Die Suche basiert auf Daten aus Proton-Proton-Kollisionen mit einer Schwerpunktsenergie von 13 TeV, die vom CMS-Detektor am LHC während des Run-2-Datennahme gesammelt wurden, entsprechend einer integrierten Luminosität von $138\,\mathrm{fb}^{-1}$. Die Analyse zielt auf Pseudoskalare mit Massen zwischen 4 und 15 GeV ab, deren stark kollimierte Zerfallsprodukte zu einem herausfordernden Endzustand mit überlappenden Fermionensignaturen führen. Es wird eine spezialisierte Strategie implementiert, die von den Standard-Rekonstruktions- und Auswahltechniken abweicht, indem Muonen und geladene Spuren zur Ereignisidentifikation verwendet werden.

Es wird keine signifikante Abweichung der Daten in Bezug auf die Erwartungen des Standardmodells in dieser Suche beobachtet. Folglich werden modellunabhängige obere Grenzen auf dem 95%-Konfidenzniveau für den Wirkungsquerschnitt mal Verzweigungsverhältnis für den 4τ-Endzustand, relativ zum SM-Higgs-Produktionsquerschnitt, festgelegt. Zusätzlich werden Ausschlussgrenzen für verschiedene Benchmark-Szenarien innerhalb des 2HDM+S-Frameworks bestimmt, wodurch der Parameterraum dieses Modells weiter eingeschränkt wird.

To my parents

Contents

1 Introduction **1**

2 Theoretical Overview **5**
2.1 The Standard Model of Particle Physics 5
2.1.1 Elementary Particles 6
2.1.2 Mathematical Formulation of the Standard Model 6
2.1.3 Quantum Electrodynamics 8
2.1.4 Quantum Chromodynamics 9
2.1.5 The Electroweak Theory 11
2.1.6 The Brout-Englert-Higgs Mechanism 14
2.1.7 The Standard Model Lagrangian 18
2.1.8 Limitations of the Standard Model 18
2.2 The Higgs Boson at the LHC 19
2.3 Higgs Beyond the Standard Model 22
2.3.1 The Two Higgs Doublet Model + Singlet 23
2.3.2 Current Constraints on 2HDM+S from Exotic Higgs Decay Searches . 29

3 Experimental Facilities **33**
3.1 The Large Hadron Collider 33
3.1.1 The Accelerator Complex 34
3.1.2 Experiments at the LHC 35
3.1.3 LHC Design and Specifications 36
3.1.4 Operation Schedule 38
3.2 The Compact Muon Solenoid Experiment 41
3.2.1 Coordinate system 41
3.2.2 Detector components 42
3.2.3 Phase 2 upgrade of the CMS detector 57

4 Event Simulation and Reconstruction in CMS **59**
4.1 Event Simulation in CMS 60
4.1.1 Hard Scattering and Parton Density Functions 60
4.1.2 Parton Showering and Hadronization 62
4.1.3 Underlying Event and Pileup 64
4.1.4 Event Generators 64
4.1.5 Detector Simulation 65

4.2 Reconstruction of Physics Objects . . . 66
4.2.1 The Particle Flow Algorithm . . . 67
4.2.2 Tracks . . . 68
4.2.3 Primary Vertices . . . 70
4.2.4 Muons . . . 71
4.2.5 Electrons and photons . . . 73
4.2.6 Jets . . . 74
4.2.7 Tau Leptons . . . 77
4.2.8 Missing Transverse Energy . . . 80

5 Search for Light Pseudoscalars in Exotic Decays of the Higgs boson 81
5.1 Motivation and Overview . . . 82
5.2 Analysis Strategy . . . 84
5.3 Data and Simulated Samples . . . 85
5.4 Triggers . . . 88
5.5 Offline Event Selection . . . 92
5.5.1 Primary Vertex Selection, Noise Filters, and HEM Veto . . . 92
5.5.2 Veto on b-tagged Jets . . . 92
5.5.3 Muon Selection . . . 94
5.5.4 Track Selection . . . 95
5.5.5 Isolation Requirements . . . 97
5.5.6 Selection of a_1 Candidates . . . 98
5.5.7 Final Selection . . . 98
5.6 Corrections to Simulations . . . 100
5.6.1 Pileup Reweighting . . . 100
5.6.2 Prefiring Weights . . . 100
5.6.3 Tracking Efficiency Corrections . . . 100
5.6.4 Muon Momentum Scale Corrections . . . 100
5.6.5 Muon Identification Efficiency Corrections . . . 100
5.6.6 Trigger Efficiency . . . 101
5.6.7 Track Isolation and Identification Efficiency . . . 104
5.6.8 Higgs-p_T Reweighting . . . 106
5.6.9 Jet Energy Corrections and b-jet Identification Efficiency . . . 106
5.7 Signal Extraction Stategy . . . 106
5.8 Signal Modeling . . . 108
5.9 Background Modeling . . . 112
5.9.1 Modeling of $f_{1D}(i)$ . . . 113
5.9.2 Modeling of $C(i, j)$ . . . 114
5.10 Systematic Uncertainties . . . 122
5.10.1 Uncertainties affecting background . . . 122
5.10.2 Uncertainties affecting signal distributions . . . 123

6 Results and Interpretation 125
6.1 Statistical Methods . . . 125
6.1.1 The Likelihood Function . . . 126
6.1.2 The Maximum Likelihood Method . . . 127

6.1.3 Hypothesis Testing . . . 128
6.2 Analysis Results . . . 130
6.2.1 Maximum-likelihood Fit to the 2D (m_1, m_2) Distribution . . . 131
6.2.2 Goodness-of-fit Test . . . 136
6.2.3 Impacts and Pulls of Nuisance Parameters . . . 138
6.2.4 Model-independent Upper Limits . . . 138
6.3 Interpretation in the context of 2HDM+S Scenarios . . . 144

7 Summary and Outlook **155**

Appendix A Alignment of the CMS Tracker during LHC Run 3 **159**
A.1 Tracker Alignment: Introduction . . . 160
A.1.1 Reference Systems . . . 160
A.1.2 Track-based Alignment Algorithms . . . 160
A.1.3 Weak Modes . . . 162
A.1.4 Alignment Strategies . . . 163
A.1.5 Alignment Validation . . . 166
A.2 Tracker Alignment during LHC Run 3 . . . 167
A.2.1 Start-up Alignments . . . 167
A.2.2 Alignment during Data Taking . . . 170
A.2.3 Alignment for Reprocessing of 2022 and 2023 data . . . 171
A.3 Summary . . . 176

Appendix B Tracker Alignment in Run 2 and Commissioning for Run 3 **177**

Appendix C Automated alignment calibration in CMS during Run 3 **185**

Appendix D Double-muon Trigger Efficiencies in 2016 and 2017 **191**

Appendix E Higgs-p_{T} Reweighting: Additional Plots **197**

Appendix F Measurement of Track Isolation and Identification Efficiency **201**
F.1 Introduction . . . 201
F.2 Event Selection . . . 202
F.3 Extraction of Scale Factors . . . 203
F.4 Dependence of Scale Factors on $\Delta R_{\mathrm{iso,trk}}$ . . . 204

Appendix G Model-independent Upper Limits **217**

Bibliography **221**

Acknowledgments **241**

Chapter 1

Introduction

"The important thing is not to stop questioning. Curiosity has its own reason for existing."

— Albert Einstein

The pursuit of understanding the universe has guided scientific exploration for several centuries. This journey has led to the development of particle physics in the twentieth century, which has facilitated significant advancements in the understanding of the fundamental components of matter. The standard model (SM) stands as a key achievement in this field, representing the current best understanding of the subatomic world. Built on the foundational principles of the special theory of relativity and quantum mechanics, the standard model is a renormalizable quantum field theory that strictly adheres to symmetry rules under gauge transformations. The model accounts for three of the four fundamental forces—strong, weak, and electromagnetic—and describes the properties and interactions of elementary particles. The SM's credibility stems from successful experimental results, including the discoveries of the W [1] and Z bosons [2], and the charm [3, 4] and top quarks [5], all of which were accurately predicted by the model.

The discovery of a scalar particle in 2012 by the ATLAS and CMS experiments at the CERN LHC, with a mass of 125 GeV and properties consistent with the standard model Higgs boson, marked a monumental milestone, confirming the last missing piece of the standard model [6–8]. This discovery validated the mechanism of electroweak symmetry breaking, which is essential for explaining how particles acquire mass. The observed properties of this particle, such as its production rate and the coupling strengths to fermions and vector bosons, demonstrate remarkable compatibility with the standard model [9, 10], thus reinforcing the validity of the SM.

Despite its monumental successes, the standard model leaves several crucial questions unanswered, prompting the development of beyond the standard model (BSM) theories. These theories seek to address the standard model's limitations, including its failure to incorporate gravity within its framework, explain neutrino masses, and account for matter-antimatter asymmetry, not to mention the absence of dark matter and dark energy. Moreover, the stan-

dard model fails to address the hierarchy problem, which concerns the need for unnaturally large adjustments to the Higgs boson's mass to protect it from large quantum fluctuations. This issue further underscores the necessity for BSM theories, which aim to offer more natural explanations without the need for such fine-tuning.

BSM theories often extend the standard model by introducing new particles, forces, dimensions, or symmetries not accounted for in the SM. Among these, supersymmetry proposes a partner particle for every particle in the standard model, potentially stabilizing the Higgs boson's mass without fine-tuning. Theories of extra dimensions suggest additional spatial dimensions beyond the observed three; some of these theories provide a framework for unifying gravity with the other fundamental forces.

Meanwhile, many BSM theories specifically involve modifications to the Higgs sector, proposing scenarios that lead to a wide array of physical Higgs states. By expanding the Higgs sector, these models propose various new interactions and decay processes. A prime example of such a theory is the Two-Higgs-Doublet Model plus a Singlet (2HDM+S) [11]. This model expands the conventional Higgs framework from the standard model by introducing a second Higgs doublet and an additional scalar singlet. These additions not only increase the potential Higgs boson states but also introduce complex dynamics, including exotic decay processes. In the 2HDM+S model, the well-known 125 GeV Higgs boson could decay into lighter scalar or pseudoscalar Higgs states. These states might then decay into standard model particles or new, invisible particles, potentially shedding light on dark matter candidates. The detection of such exotic decays of the 125 GeV Higgs boson would act as a direct probe of new physics. Experimental searches for these unusual decay channels are, therefore, pivotal.

Exploring the possibility of the 125 GeV Higgs boson decaying into lighter pseudoscalar particles forms the core motivation behind the research in this thesis. The thesis presents a search for light pseudoscalars produced in the decays of the 125 GeV Higgs boson, particularly examining their subsequent decay into pairs of tau leptons which result in a four-tau final state. To enhance the sensitivity of the search, the study also leverages final states consisting of two muons and two taus.

The analysis utilizes proton-proton collision data corresponding to an integrated luminosity of $138\,\mathrm{fb}^{-1}$, recorded by the CMS detector at the LHC during the Run 2 data taking period, at a center-of-mass energy of 13 TeV. This study targets masses of the light pseudoscalar ranging from 4 to 15 GeV. Due to the considerable mass difference between these light pseudoscalars and the 125 GeV Higgs boson, the light pseudoscalars are produced with a high degree of boost, resulting in collimated decay products. Such characteristics necessitate a specialized event selection strategy designed to effectively identify these decays. The final state, though challenging, is worth exploring as the leptonic decay of the pseudoscalar to taus is highly favored in a considerable portion of the phase space of various BSM models.

The thesis is structured as follows: Chapter 2 lays the groundwork for the analysis by providing a comprehensive theoretical overview of the standard model of particle physics, emphasizing its limitations, and segueing into a discussion on various beyond the standard model theories that predict exotic decays of the 125 GeV Higgs boson. Chapter 3 describes the experimental facilities, focusing on the design and workings of the LHC and the CMS detector. The quality of data used for physics analyses depends critically on rigorous calibrations of

the detector subsystems. Among these, the alignment of the CMS tracker is essential for accurate particle tracking and vertex reconstruction. The author of this thesis has made substantial contributions to the tracker alignment process during the duration of their PhD. These efforts are briefly mentioned in Chapter 3 and discussed in detail in the Appendix A. This work has resulted in several publications, listed in Refs. [12–17]. Chapter 4 then explores the methodologies used in event simulation and the reconstruction of physics objects within CMS, essential for carrying out a physics analysis.

Chapter 5 discusses the core work carried out by the author during their PhD. This chapter describes the search for light pseudoscalars produced in decays of the 125 GeV Higgs boson, in final states involving either four tau leptons or two muons and two tau leptons. The chapter details the comprehensive analysis of data from 2016, 2017, and 2018. It starts with a brief overview of the search, describing the special strategy employed to identify highly Lorentz-boosted tau lepton pairs with overlapping decays. It also covers the approach to signal modeling and background estimation using data-driven techniques. Chapter 6 presents the results for each year of data analyzed, as well as for the full Run 2 combination, detailing the statistical methods used. It provides the model-independent upper limits set by the analysis on the cross section times branching ratio for the 4τ final state, relative to the standard model Higgs production cross section. The chapter also highlights the model-dependent exclusion limits derived for various benchmark 2HDM+S scenarios. It is important to emphasize that the findings from this analysis have been released publically in a CMS physics analysis summary [18] and are currently in the process of being finalized for journal publication. Lastly, Chapter 7 summarizes the key findings and provides an outlook on potential future research directions.

Chapter 2

Theoretical Overview

Contents

2.1 The Standard Model of Particle Physics 5
2.1.1 Elementary Particles 6
2.1.2 Mathematical Formulation of the Standard Model 6
2.1.3 Quantum Electrodynamics 8
2.1.4 Quantum Chromodynamics 9
2.1.5 The Electroweak Theory 11
2.1.6 The Brout-Englert-Higgs Mechanism 14
2.1.7 The Standard Model Lagrangian 18
2.1.8 Limitations of the Standard Model 18
2.2 The Higgs Boson at the LHC 19
2.3 Higgs Beyond the Standard Model 22
2.3.1 The Two Higgs Doublet Model + Singlet 23
2.3.2 Current Constraints on 2HDM+S from Exotic Higgs Decay Searches 29

This chapter provides an overview of the standard model of particle physics, highlighting its fundamental features. It then discusses its limitations, which motivate the need for beyond-standard-model theories. Building on this need for new physics, the discussion transitions to extended Higgs sectors, specifically focusing on the Two-Higgs-Doublet Model with an additional singlet (2HDM+S). This model offers a rich phenomenology for exotic decays of the 125 GeV Higgs boson and sets the theoretical groundwork for the analysis presented in this thesis.

2.1 The Standard Model of Particle Physics

The standard model (SM) of particle physics is a relativistic quantum field theory (QFT) based on the principles of gauge invariance. Developed and refined in the latter half of the 20th century, the SM successfully accounts for three of the four fundamental forces—strong, weak, and electromagnetic—and describes the properties and interactions of elementary particles. The following subsections offer a concise overview of the key aspects of the model.

2.1.1 Elementary Particles

The spectrum of elementary particles in the standard model is categorized into bosons and fermions based on their intrinsic angular momentum. Fermions, which have half-integer spins, constitute matter, while bosons, with integer spins, play the role of force carriers.

Fermions are further categorized into two groups: quarks and leptons. Quarks exist in six *flavors*: up (u), down (d), charm (c), strange (s), top (t), and bottom (b). The u, c, and t-quarks carry an electric charge of $+2/3\,e$ and are called the up-type quarks. On the other hand, the d, s, and b-quarks, also called the down-type quarks, possess an electric charge of $-1/3\,e$. In addition to the electromagnetic charge, each quark flavor possesses a color charge—red, green, or blue—enabling them to participate in strong interactions. The masses of quarks vary significantly, from a few MeV for the up and down quarks to 173 GeV for the top quark.

Similarly, there are six leptons: the electrically-charged electron (e), muon (μ) and tau lepton (τ), accompanied by a corresponding neutral particle, the electron neutrino (ν_e), muon neutrino (ν_μ) and tau neutrino (ν_τ). All charged leptons carry an electric charge equal to $-e$. Unlike quarks, leptons do not possess a color charge and, therefore, do not participate in strong interactions. The masses of charged leptons range from 511 keV for electrons to 1.777 GeV for tau leptons. Within the SM framework, neutrinos are postulated to be massless.

Quarks and leptons are organized into three generations, with particles from different generations sharing the same properties but having different masses. Each quark generation comprises an up-type and a down-type quark, whereas a lepton generation consists of a charged lepton and its corresponding neutrino. Furthermore, every fermion has a corresponding antiparticle, which has the same mass and spin but carries the opposite quantum numbers.

Besides fermions, the SM also includes 12 gauge or vector bosons, which are spin-1 particles responsible for mediating the fundamental forces. The massless photon mediates electromagnetic interactions among electrically charged particles. The $W^{\pm}$ and Z bosons, with masses of 80 GeV and 91 GeV, respectively, govern the weak force. The $W^{\pm}$ bosons mediate the weak charged-current interactions that can alter the flavor of fermions, whereas the Z boson mediates weak neutral-current interactions that do not alter particle flavors. Eight gluons (g) are responsible for binding quarks together through the strong force, forming color-neutral composite states known as hadrons (e.g., protons).

The standard model also accommodates the Higgs boson, which is a massive scalar boson with zero spin. This particle is responsible for conferring mass to itself and the other massive particles with the SM.

An overview of the particle content of the standard model is provided in Fig. 2.1.

2.1.2 Mathematical Formulation of the Standard Model

As mentioned, the standard model is a quantum field theory. This means that it treats particles as excitations of underlying quantum fields characterized by their quantum numbers, such as charge and spin. The quantum fields are dynamic entities that permeate all of space.

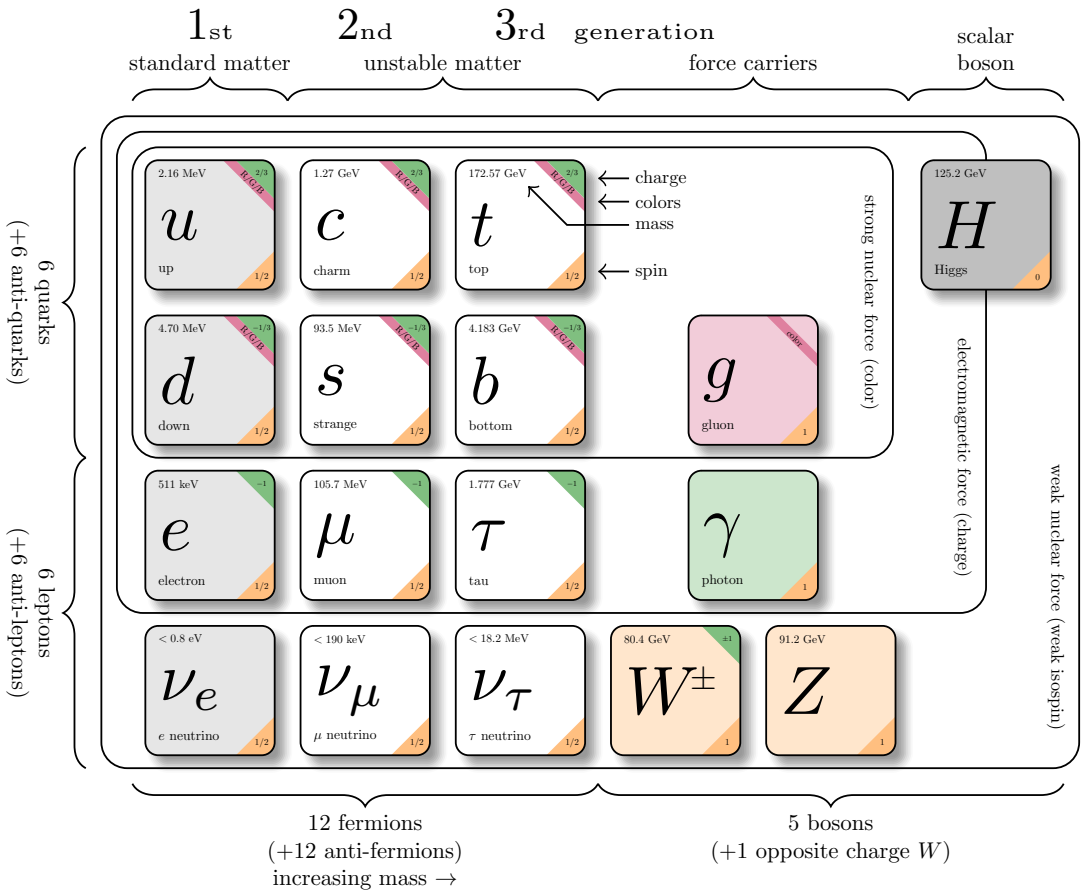

Figure 2.1: Overview of the fundamental particles in the standard model, displaying the three generations of fermions (quarks and leptons), gauge bosons, and the Higgs boson. Also shown are the mass, charge, and spin properties of each particle. The given figure is adapted from Ref. [19], with values taken from Ref. [20].

To describe the dynamics of these fields, the Lagrangian formalism is employed, where each field and its interactions are represented through a Lagrangian density $\mathcal{L}$, that includes terms for kinetic energies, mass, and field interactions.

The mathematical formulation of the standard model is constructed to respect fundamental symmetries, particularly gauge symmetries. These symmetries dictate the form of the Lagrangian and, through Noether's theorem [21], lead to the conservation of quantities such as electric charge, weak isospin, and color charge.

Building on these principles, the SM Lagrangian is formulated by incorporating the Lagrangian densities corresponding to each fundamental interaction. The construction of these individual components and their contributions to the complete theory are discussed in subsequent sections.

2.1.3 Quantum Electrodynamics

Quantum Electrodynamics (QED) is the quantum field theory that describes electromagnetic interactions and is the first successful theory formulated based on the principle of gauge invariance [22].

The QED Lagrangian can be derived by imposing local $U(1)$ gauge invariance on the Dirac Lagrangian for free fermions [23]. In natural units ($\hbar = c = 1$), the Dirac Lagrangian is formulated as:

$$\mathcal{L}_{\text{Dirac}} = i\overline{\psi}\gamma^{\mu}\partial_{\mu}\psi - m\overline{\psi}\psi, \tag{2.1}$$

where ψ represents the fermionic field, $\overline{\psi} = \psi^{\dagger}\gamma^{0}$ is the adjoint of the fermionic field, and γ^{μ} are the Dirac matrices. The first term here represents the fermion kinetic term, while the second term represents the mass term.

Under $U(1)$ gauge transformations, the fermionic field transforms as $\psi \to \psi' = e^{i\alpha(x)}\psi$, where $\alpha(x)$ is a space-time dependent phase. This transformation requires that the standard derivative ∂_{μ} in the Lagrangian be replaced with the covariant derivative D_{μ} to ensure gauge invariance. In QED, D_{μ} is defined as:

$$D_{\mu} = \partial_{\mu} + iqA_{\mu}, \tag{2.2}$$

where q is a constant. A_{μ} is a gauge field that transforms as $A_{\mu} \to A_{\mu} + \partial_{\mu}\alpha$.

Substituting this covariant derivative into the Dirac Lagrangian yields the gauge-invariant QED Lagrangian:

$$\begin{aligned}\mathcal{L}_{\text{QED}} &= \overline{\psi}(i\gamma^{\mu}D_{\mu} - m)\psi \\ &= i\overline{\psi}\gamma^{\mu}\partial_{\mu}\psi - m\overline{\psi}\psi - q\overline{\psi}\gamma^{\mu}\psi A_{\mu}.\end{aligned} \tag{2.3}$$

This introduces a new interaction term, $-q\overline{\psi}\gamma^{\mu}A_{\mu}\overline{\psi}$, which describes the exchange of the gauge field A_{μ} between fermions. Consequently, A_{μ} can be identified as the photon field.

To allow A_{μ} to propagate, an additional kinetic term is defined for the gauge field:

$$\mathcal{L}_{\text{EM}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \tag{2.4}$$

where $F_{\mu\nu}$, the field strength tensor, is defined as $F_{\mu\nu} = \partial_{\mu}A_{\nu} - \partial_{\nu}A_{\mu}$. No mass term corresponding to A_{μ} is introduced as it would otherwise break the gauge invariance of the theory; consequently, the photon remains massless, consistent with experimental observations [20].

The complete gauge-invariant QED Lagrangian, incorporating both the fermionic and gauge field components, is expressed as:

$$\mathcal{L}_{\text{QED}} = i\overline{\psi}\gamma^{\mu}\partial_{\mu}\psi - m\overline{\psi}\psi - q\overline{\psi}\gamma^{\mu}\psi A_{\mu} - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}. \tag{2.5}$$

The invariance of the Lagrangian under local $U(1)$ transformations leads to the conservation of electric charge, with $q = -e$ corresponding to the charge of the electron.

2.1.4 Quantum Chromodynamics

Quantum Chromodynamics (QCD) is the fundamental theory of strong interaction, responsible for binding quarks and gluons into hadrons [24]. Experimental evidence regarding the composition of hadrons led to the inference that quarks possess an additional quantum number compared to leptons, called the *color charge*. In QCD, the color charge manifests in three types, commonly referred to as red, green, and blue. This leads to the description of strong interactions through the non-abelian $SU(3)_C$ symmetry group, with C representing the conserved color charge [23].

Under this group, quark fields ψ transform as follows:

$$\psi \to \psi' = U\psi, \tag{2.6}$$

where U is an $SU(3)$ matrix dependent on space-time coordinates. These matrices are unitary with a determinant of one, characteristic of the $SU(3)$ group. The matrix can be mathematically expressed in exponential form using the generators T^a of $SU(3)$:

$$U = \exp(i\theta^a T^a), \quad T^a = \frac{\lambda^a}{2}, \quad a = 1, 2, \ldots, 8 \tag{2.7}$$

where T^a are the generators expressed in terms of the Gell-Mann matrices λ^a and θ^a are the parameters of the transformation.

To ensure the gauge invariance of the QCD Lagrangian under these local $SU(3)$ transformations, the covariant derivative is defined as:

$$D_\mu = \partial_\mu + ig_s T^a G^a_\mu, \tag{2.8}$$

where g_s is a constant, and G^a_μ correspond to eight gauge fields that transform according to:

$$G^a_\mu \to G'^a_\mu = G^a_\mu - \frac{1}{g_s}\partial_\mu \theta^a - f^{abc}\theta^b G^c_\mu. \tag{2.9}$$

Here f^{abc} are the structure constants of the SU(3) group, defining the non-commutative properties of the generators[1].

The full QCD Lagrangian is thus given by:

[1] The generators of the SU(3) group follow the commutation rule $[T^a, T^b] = if^{abc}T^c$.

$$\begin{aligned}\mathcal{L}_{\text{QCD}} &= \sum_f \overline{\psi}_f (i\gamma^\mu D_\mu - m)\psi_f - \frac{1}{4} G^a_{\mu\nu} G^{a\mu\nu} \\ &= \sum_f \overline{\psi}_f (i\gamma^\mu \partial_\mu - m)\psi_f - \sum_f g_s \overline{\psi}_f \gamma^\mu G^a_\mu T^a \psi_f - \frac{1}{4} G^a_{\mu\nu} G^{a\mu\nu}.\end{aligned} \tag{2.10}$$

In this expression:

- The first term, $\overline{\psi}_f(i\gamma^\mu \partial_\mu - m)\psi_f$, describes the kinematics and mass of the quarks of flavor f.
- The second term, $-g_s \overline{\psi}_f \gamma^\mu T^a G^a_\mu \psi_f$, represents the interaction of the quark field with the gauge fields. The gauge fields can thus be identified as the eight gluon fields mediating the strong interaction. The constant g_s represents the strong coupling constant.
- The gluon kinetic term, $\frac{1}{4} G^a_{\mu\nu} G^{a\mu\nu}$, encapsulates the dynamics of the gluon fields.

The gluon field strength tensor, crucial for understanding the dynamics of gluons, is defined as:

$$G^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu + g_s f^{abc} G^b_\mu G^c_\nu. \tag{2.11}$$

While the first two terms of $G^a_{\mu\nu}$ are akin to the electromagnetic field strength tensor in QED, the third term, $g_s f^{abc} G^b_\mu G^c_\nu$ is unique to the non-Abelian nature of QCD. This term introduces trilinear and quartic gluon couplings when the tensor is substituted back into the Lagrangian, highlighting the self-interactions among gluons. These interactions are illustrated in Fig. 2.2, which lead to the very complex kinematics and rich phenomenology of the strong interactions.

Figure 2.2: Feynman diagrams showing gluon self-interactions: (left) three-point coupling, (right) four-point coupling.

A profound implication of the gluon self-interaction in QCD is the variation of the strong coupling constant α_s with the energy scale of the interaction. The behavior is described as:

$$\alpha_s \equiv \frac{g_s^2}{4\pi} = \frac{12\pi}{(33 - 2n_f)\log(Q^2/\Lambda^2)}, \tag{2.12}$$

where n_f is the number of active quark flavors at the energy scale Q^2, and Λ represents the QCD scale typically around 100-300 MeV, varying based on n_f [20]. As Q^2 increases, α_s decreases. At very high energy scales, the strong force is virtually absent, and quarks and gluons behave almost as free particles. This phenomenon is called *asymptotic freedom* [25,

26]. In the high-energy regime, perturbative techniques can be used to study the strong interactions. In contrast, at low Q^2 close to Λ, α_s rises sharply, leading to strong coupling. This strong interaction leads to *color confinement*, a process that binds quarks and gluons into composite, color-neutral particles called hadrons [27]. Consequently, gluons and quarks are never observed as isolated entities outside of hadrons. At these lower energies, the increase in α_s renders perturbative methods ineffective, requiring non-perturbative approaches to study these strong interaction phenomena.

2.1.5 The Electroweak Theory

The electroweak theory, formulated by Glashow [28], Weinberg [29], and Salam [30], unifies the electromagnetic and weak nuclear forces within a single theoretical framework governed by the gauge group $SU(2)_L \times U(1)_Y$.

The $SU(2)_L$ and $U(1)_Y$ symmetries are generated by the weak isospin operators $\vec{T}$ and the hypercharge operator Y, respectively. The former is expressed in terms of Pauli matrices as:

$$\vec{T} = \frac{\vec{\sigma}}{2}. \tag{2.13}$$

These generators satisfy non-abelian commutation relations, reflecting the complex structure of the weak interactions:

$$[T_a, T_b] = i\epsilon_{abc} T_c. \tag{2.14}$$

Here ϵ_{ijk} is the Levi-Civita tensor. In contrast, the hypercharge operator Y commutes with all $SU(2)_L$ generators as:

$$[T_a, Y] = 0, \tag{2.15}$$

reflecting the independent action of the $U(1)_Y$ symmetry. The hypercharge Y and the third component of the isospin I_3, are related to the electromagnetic charge Q, as:

$$Y = 2(Q - I_3). \tag{2.16}$$

The electroweak theory is particularly notable for its chiral nature, which differentiates between left-handed and right-handed fermions. Chirality or handedness refers to the orientation of a particle's spin relative to its momentum. Left-handed fermions have spins pointing opposite to their direction of motion, while right-handed fermions have spins aligned with their motion. Chirality is mathematically represented using projection operators, P_L and P_R, defined as:

$$P_L = \frac{1}{2}(1 - \gamma^5), \quad P_R = \frac{1}{2}(1 + \gamma^5), \tag{2.17}$$

where $\gamma^5 = -i\gamma^0\gamma^1\gamma^2\gamma^3$ is the fifth gamma matrix. These operators decompose any fermion field ψ into its left-handed (ψ_L) and right-handed (ψ_R) components as:

$$\begin{aligned} \psi_L &= P_L\psi = \frac{1}{2}(1-\gamma^5)\psi, \\ \psi_R &= P_R\psi = \frac{1}{2}(1+\gamma^5)\psi \end{aligned} \tag{2.18}$$

In the electroweak theory, left-handed leptons and quarks are grouped into doublets under $SU(2)_L$ as:

$$\psi_L^\ell = \begin{pmatrix} \nu_L \\ \ell_L \end{pmatrix}, \quad \psi_L^q = \begin{pmatrix} \mathrm{u}_L \\ \mathrm{d}_L \end{pmatrix} \tag{2.19}$$

where ℓ, ν, u, and d represent leptons, corresponding lepton neutrinos, up-type quarks and down-type quarks, respectively [31]. On the other hand, right-handed fermions are treated as singlets:

$$\ell_R, \quad \mathrm{u}_R, \quad \mathrm{d}_R, \tag{2.20}$$

interacting only through the $U(1)_Y$ mediated processes. This distinction in the treatment of left- and right-handed fermions is crucial to accommodate the experimentally observed parity violation in weak interactions, as first shown in the 1957 Cobalt-60 experiment by Wu [32]. These findings demonstrated that the weak force only interacts with left-handed fermions, which the electroweak theory effectively integrates into its framework.

Turning to the gauge transformations, a generic left-handed field ψ_L and right-handed field ψ_R transform under $SU(2)_L \times U(1)_Y$ as:

$$\begin{aligned} \psi_L \to \psi_L' &= \exp\left(i\vec{\alpha}\cdot\vec{T} + i\beta Y_L\right)\psi_L, \\ \psi_R \to \psi_R' &= \exp\left(i\beta Y_R\right)\psi_R \end{aligned} \tag{2.21}$$

where $\vec{\alpha}$ and β represent the parameters of the transformations in $SU(2)_L$ and $U(1)_Y$, respectively. Y_L and Y_R denote the hypercharges for the left-handed doublets and right-handed singlets, respectively.

Following a similar approach as for QED and QCD, the electroweak Lagrangian can be derived starting from the free Lagrangian for the fermion fields and requiring it to be invariant under the $SU(2)_L \times U(1)_Y$ transformations. The Lorentz derivative ∂_μ for the fermion field is replaced by the covariant derivatives $D_{\mu,L}$ and $D_{\mu,R}$ for left-handed doublets and the right-handed singlets respectively:

$$\begin{aligned} D_{\mu,L} &= \partial_\mu + ig\vec{T}\vec{W}_\mu + ig'\frac{Y_L}{2}B_\mu, \\ D_{\mu,R} &= \partial_\mu + ig'\frac{Y_R}{2}B_\mu. \end{aligned} \tag{2.22}$$

The requirement of the $SU(2)_L \times U(1)_Y$ gauge invariance naturally leads to the introduction of $3 + 1 = 4$ gauge fields—$\vec{W} = \{W_1, W_2, W_3\}$ associated with the weak isospin, and B_μ associated with the hypercharge. Here, g and g' are the $SU(2)_L$ and $U(1)_Y$ coupling constants, respectively.

Assuming massless fermions, the electroweak Lagrangian, including the kinetic terms for the gauge bosons, can be written as:

$$\begin{aligned}\mathcal{L}_{\text{EW}} = &\sum_f \left(\overline{\psi}^{\ell}_{L,f}\gamma^\mu D_{\mu,L}\psi^{\ell}_{L,f} + \overline{\ell}_{R,f}\gamma^\mu D_{\mu,R}\ell_{R,f}\right) \\ &+ \sum_f \left(\overline{\psi}^{q}_{L,f}\gamma^\mu D_{\mu,L}\psi^{q}_{L,f} + \overline{u}_{R,f}\gamma^\mu D_{\mu,R}u_{R,f} + \overline{d}_{R,f}\gamma^\mu D_{\mu,R}d_{R,f}\right) \\ &- \frac{1}{4}W^a_{\mu\nu}W^{a\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu},\end{aligned} \tag{2.23}$$

where f runs over the three families of quarks and leptons. The field strength tensors, $B_{\mu\nu}$ and $W^a_{\mu\nu}$, are defined as:

$$\begin{aligned}B_{\mu\nu} &= \partial_\mu B_\nu - \partial_\nu B_\mu, \\ W^a_{\mu\nu} &= \partial_\mu W^a_\nu - \partial_\nu W^a_\mu - g\epsilon^{abc}W^b_\mu W^c_\nu.\end{aligned} \tag{2.24}$$

The fields W^a_μ and B_μ mix as follows to form the fields of the physical electroweak gauge bosons:

$$\begin{aligned}W^{\pm}_\mu &= \frac{1}{\sqrt{2}}(W^1_\mu \pm iW^2_\mu), \\ Z_\mu &= \frac{gW^3_\mu - g'B_\mu}{\sqrt{g^2 + g'^2}} = W^3_\mu cos\,\theta_W - B_\mu sin\,\theta_W, \\ A_\mu &= \frac{g'W^3_\mu + gB_\mu}{\sqrt{g^2 + g'^2}} = W^3_\mu sin\,\theta_W + B_\mu cos\,\theta_W.\end{aligned} \tag{2.25}$$

Here, θ_W, also known as the Weinberg angle, denotes the mixing angle in the neutral weak boson sector. It is expressed in terms of coupling constants as:

$$cos\,\theta_W = \frac{g}{\sqrt{g^2 + g'^2}} \tag{2.26}$$

The value of the Weinberg angle is not predicted within the SM but is determined experimentally to be $\sin^2\theta_W = 0.23129 \pm 0.00002$, measured at the Z boson's mass [20].

The unification of electromagnetic and weak forces into the electroweak theory represents a landmark achievement in particle physics. However, it is affected by a significant limitation. In the formulation described above, the $W^{\pm}$ and Z bosons appear massless, contradicting experimental observations which show these bosons are indeed massive. Introducing a gauge mass term $MW^\mu W_\mu$ directly into the electroweak Lagrangian would violate gauge invariance.

Similarly, adding fermion mass terms of the form $-m^2\psi\overline{\psi}$, which involve mixed chirality, would not align with the gauge transformation rules.

The requirement to maintain gauge invariance while introducing mass naturally leads to the need for a symmetry-breaking mechanism. The electroweak symmetry $SU(2)_L \times U(1)_Y$ must be broken down to the charge symmetry of electromagnetism, $U(1)_{\mathrm{EM}}$, to align the theory with experimental observations and allow both gauge bosons and fermions to acquire mass consistently. The subsequent section elaborates on the mechanism responsible for this essential symmetry breaking.

2.1.6 The Brout-Englert-Higgs Mechanism

The Brout-Englert-Higgs (BEH) mechanism, developed independently by Brout and Englert and by Higgs, addresses the generation of gauge boson masses and the presence of fermion masses in the electroweak theory [33–37]. This mechanism introduces a scalar field that spontaneously breaks the electroweak symmetry. Spontaneous symmetry breaking (SSB) occurs when the symmetric state of a system becomes unstable, and the system settles into one of multiple possible asymmetric states under small perturbations, even though its governing laws remain symmetric. This concept is well-established in physics, with applications ranging from the formation of crystals in condensed matter physics to the evolution of structure in the early universe in cosmology.

In the standard model, the BEH mechanism is implemented using a complex scalar field, Φ, commonly referred to as the Higgs field. It forms a doublet under $SU(2)_L$ and is mathematically expressed as:

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}. \tag{2.27}$$

The Higgs field is required to be a scalar to maintain the isotropy and homogeneity of space, avoiding directional biases that could break fundamental symmetries. A constant vacuum expectation value (VEV) is crucial to ensure uniform particle mass distribution. Additionally, it is required to have a hypercharge of $Y = 1$ to allow the field to remain uncharged under $U(1)_{\mathrm{EM}}$ after electroweak symmetry breaking.

The corresponding Lagrangian can be written as:

$$\mathcal{L}_{BEH} = (D^\mu)^\dagger (D_\mu)\Phi - V(\Phi) \tag{2.28}$$

where the covariant derivate D_μ corresponds to the one introduced in Eq. 2.22. The potential $V(\Phi)$ is defined as:

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda(\Phi^\dagger \Phi)^2 = \mu^2 |\Phi|^2 + \lambda |\Phi|^4 \tag{2.29}$$

Here, λ is required to be strictly positive in order to ensure the stability of the Higgs field's vacuum, preventing the potential from descending into negative values for large Φ. By requiring $\mu^2 < 0$, the potential adopts a degenerate minimum, characterized by the familiar

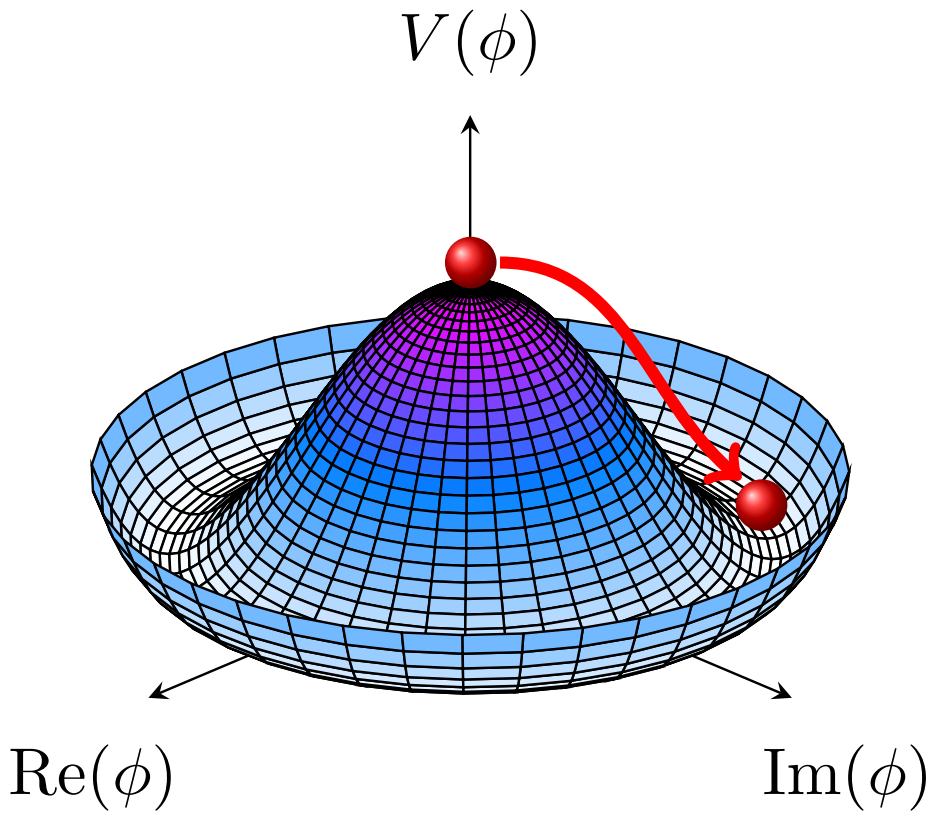

Figure 2.3: Illustration of the functional form of the Higgs potential. The potential presents a local maximum for $\Phi = 0$ and degenerate minima at $\frac{-\mu^2}{2\lambda}$. The red sphere depicts the process of passing from an unstable local maximum to the continuum of ground states.

sombrero shape as depicted in Fig. 2.3. This configuration results in a vacuum ground state that possesses infinitely many equivalent solutions, each represented by:

$$|\Phi|^2 = \frac{-\mu^2}{2\lambda}. \tag{2.30}$$

The $SU(2)_L \times U(1)_Y$ symmetry is broken by selecting a specific ground state from the continuum of degenerate vacua. The conventional choice sets $\phi^+_{vac} = 0$ and $\phi^0_{vac} = v$, establishing the ground state Φ_0 as:

$$\Phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \tag{2.31}$$

This choice preserves the $U(1)_{\mathrm{EM}}$ symmetry, thereby ensuring the conservation of electric charge. The Lagrangian specified in Eq. 2.28 can be perturbatively expanded around the chosen ground state Φ_0 by reparameterizing the scalar field as:

$$\Phi = \frac{1}{\sqrt{2}} \exp\left(\frac{i}{v}\sigma_j \theta^j\right) \begin{pmatrix} 0 \\ v + \mathrm{H} \end{pmatrix} \tag{2.32}$$

Here, H represents the physical Higgs boson, and θ^j $(j = 1, 2, 3)$ are massless scalar fields corresponding to Goldstone bosons. These Goldstone bosons arise as a consequence of Goldstone's theorem, which states that the spontaneous breaking of a continuous symmetry results in the appearance of massless bosons in quantities equal to the number of broken symmetry generators [38]. With the breaking of $SU(2)_L \times U(1)_Y$ to $U(1)_{\mathrm{EM}}$, three of the four generators are broken, giving rise to three Goldstone bosons.

Since the massless Goldstone bosons do not correspond to physical particles, they must be absorbed into the theory. This is accomplished by applying a local $SU(2)_L$ gauge transformation, known as the unitary gauge, to the Higgs field doublet. This transformation is expressed as:

$$\Phi \to \exp\left(-\frac{i}{v}\sigma_j \theta^j\right) \Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + \mathrm{H} \end{pmatrix}. \tag{2.33}$$

The BEH Lagrangian can then be rewritten as:

$$\begin{aligned} \mathcal{L}_{BEH} = & \frac{1}{2}\partial_\mu H \partial^\mu H - \lambda v^2 H^2 - \lambda v H^3 - \frac{\lambda}{4} H^4 \\ & + \frac{g^2 v^2}{8}\left(W_\mu^1 W^{1\mu} + W_\mu^2 W^{2\mu}\right) + \frac{v^2}{8}\left(g W_\mu^3 - g' B_\mu\right)^2 \\ & + \frac{g^2 v H}{4}\left(W_\mu^1 W^{1\mu} + W_\mu^2 W^{2\mu}\right) + \frac{v H}{4}\left(g W_\mu^3 - g' B_\mu\right)^2 \\ & + \frac{g^2 H^2}{8}\left(W_\mu^1 W^{1\mu} + W_\mu^2 W^{2\mu}\right) + \frac{H^2}{8}\left(g W_\mu^3 - g' B_\mu\right)^2. \end{aligned} \tag{2.34}$$

Taking into consideration the relations given in Eq. 2.25, three explicit mass terms for the gauge bosons can be identified from the Lagrangian above:

$$\begin{aligned} m_{\mathrm{H}} &= \sqrt{2\lambda v} = \sqrt{-2\mu^2}, \\ m_W &= \frac{g v}{2}, \\ m_Z &= \frac{v}{2}\sqrt{g^2 + g'^2}. \end{aligned} \tag{2.35}$$

Here, m_W and m_Z satisfy the relation:

$$m_Z = \frac{m_W}{\cos\theta_W}. \tag{2.36}$$

No explicit mass term appears for the photon field $A_\mu = g W_\mu^3 + g' B_\mu$, ensuring it remains massless, as expected from the unbroken $U(1)_{\mathrm{EM}}$ symmetry.

The mass of the Higgs boson, m_{H}, is a free parameter of the theory and is directly dependent on the vacuum expectation value and the self-coupling constant λ. The vacuum expectation value itself can be determined from the experimentally measured Fermi constant G_F [23] as follows:

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{m_W^2} = \frac{1}{2v^2} \implies v = \frac{1}{(\sqrt{2}G_F)^{1/2}} \approx 246\,\mathrm{GeV}. \tag{2.37}$$

While the BEH mechanism and electroweak symmetry breaking provide masses to the weak bosons, they do not directly affect fermion masses. To address this, an additional mechanism—the Yukawa interaction—is introduced [29]. The Yukawa Lagrangian is given by:

$$\mathcal{L}_{\mathrm{Yukawa}} = \sum_{ij}(-Y_u^{ij}\bar{\psi}^{q^i}{}_L\tilde{\Phi}u_R^j - Y_d^{ij}\bar{\psi}^{q^i}{}_L\Phi d_R^j - Y_\ell^{ij}\bar{\psi}^{\ell^i}{}_L\Phi\ell_R^j + \mathrm{h.c.}), \tag{2.38}$$

where Y_u^{ij}, Y_d^{ij}, and Y_ℓ^{ij} are the Yukawa coupling matrices for up-type quarks, down-type quarks, and charged leptons, respectively; i, j are generation indices; and h.c. refers to the hermitian conjugate of the terms. Here, $\tilde{\Phi}$, which corresponds to the conjugate Higgs doublet:

$$\tilde{\Phi} = i\sigma^2\Phi^* = \begin{pmatrix}(\phi^0)^*\\ -\phi^-\end{pmatrix} \xrightarrow{SSB} \frac{1}{\sqrt{2}}\begin{pmatrix}v+H\\ 0\end{pmatrix} \tag{2.39}$$

is introduced to generate masses of the up-type quarks.

Following electroweak symmetry breaking, the Lagrangian acquires the form:

$$\begin{aligned}\mathcal{L}_{\mathrm{Yukawa}} = &-\frac{v}{\sqrt{2}}\sum_{ij}\left(Y_u^{ij}\bar{u}_L^i u_R^j + Y_d^{ij}\bar{d}_L^i d_R^j + Y_\ell^{ij}\bar{\ell}_L^i\ell_R^j\right)\\ &-\frac{H}{\sqrt{2}}\sum_{ij}\left(Y_u^{ij}\bar{u}_L^i u_R^j + Y_d^{ij}\bar{d}_L^i d_R^j + Y_\ell^{ij}\bar{\ell}_L^i\ell_R^j\right) + \mathrm{h.c.}\end{aligned} \tag{2.40}$$

Here, the first line corresponds to the fermion mass terms, and the second line corresponds to the interaction between the fermion and the Higgs boson.

The masses M_f^{ij} can be inferred as :

$$M_f^{ij} = \frac{v}{\sqrt{2}}Y^{ij}, \tag{2.41}$$

However, the mass matrix M_f is generally not diagonal in the weak interaction basis, meaning the mass eigenstates do not align with the weak eigenstates. To obtain the physical masses, M_f is diagonalized via a bi-unitary transformation:

$$M_f^{\mathrm{diag}} = V_L^\dagger M_f V_R, \tag{2.42}$$

where V_L and V_R are unitary matrices. This diagonalization results in flavor mixing among quark generations, giving rise to the Cabibbo-Kobayashi-Maskawa (CKM) matrix, which governs weak interaction transitions between quark flavors[2] [40, 41]. In the lepton sector, since neutrinos are assumed to be massless, no analogous mixing occurs.

Furthermore, it is evident from Eq. 2.40 that the Yukawa interaction allows the Higgs boson to couple to fermions with a strength proportional to their mass, explaining why heavier fermions interact more strongly with the Higgs field.

2.1.7 The Standard Model Lagrangian

As mentioned earlier in Sec. 2.1.2, the SM Lagrangian can be constructed by combining the Lagrangian for the individual interactions. Therefore, considering all the components discussed in the previous sections, the SM Lagrangian can be formulated from a $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetry as:

$$\begin{aligned}
\mathcal{L}_{\text{SM}} = & -\frac{1}{4} W^a_{\mu\nu} W^{a\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} G^a_{\mu\nu} G^{a\mu\nu} \\
& + \sum_f \left(\overline{\psi}^{\ell}_{L,f} \gamma^\mu D_\mu \psi^{\ell}_{L,f} + \overline{\ell}_{R,f} \gamma^\mu D_\mu \ell_{R,f} \right) \\
& + \sum_f \left(\overline{\psi}^{q}_{L,f} \gamma^\mu D_\mu \psi^{q}_{L,f} + \overline{u}_{R,f} \gamma^\mu D_\mu u_{R,f} + \overline{d}_{R,f} \gamma^\mu D_\mu d_{R,f} \right) \\
& + (D^\mu \Phi)^\dagger (D_\mu \Phi) - V(\Phi) \\
& + \sum_{ij} (-Y_u^{ij} \bar{\psi}^{q^i}_L \tilde{\Phi} u^j_R - Y_d^{ij} \bar{\psi}^{q^i}_L \Phi d^j_R - Y_\ell^{ij} \bar{\psi}^{\ell^i}_L \Phi \ell^j_R + \text{h.c.}).
\end{aligned} \tag{2.43}$$

The covariant derivative $D_\mu = \partial_\mu + ig\vec{T}\vec{W}_\mu + ig' \frac{Y_L}{2} B_\mu + ig_s T^a G^a_\mu$ ensures the invariance under local gauge transformations associated with each symmetry group. This formulation of the Standard Model Lagrangian thus elegantly condenses the current understanding of the universe into one compact equation.

2.1.8 Limitations of the Standard Model

With the discovery of the Higgs boson in 2012 [6–8], the last missing piece of the standard model was put in place, reaffirming it as the most robust framework to date for describing fundamental particles and their interactions. Despite its successes, the SM is widely acknowledged as not being a complete theory, as it does not account for several critical phenomena.

The SM does not address the existence of dark matter and dark energy, which together are estimated to constitute about 95% of the universe's total mass-energy content [20, 42].

[2]The CKM matrix contains a complex phase that introduces a violation of charge-parity (CP) symmetry. This symmetry involves charge conjugation (C), which swaps particles with their antiparticles, and parity (P), which inverts spatial coordinates like a mirror reflection. Despite CP symmetry suggesting identical interactions for particles and antiparticles, experiments like the 1964 Cronin-Fitch experiment with kaon decays confirmed its violation in weak interactions [39]. This asymmetry in decay rates between particles and their antiparticles is crucial for explaining the observed matter-antimatter imbalance in the universe.

Additionally, it does not incorporate gravity, one of the four fundamental forces, nor does it provide a mechanism to explain the observed non-zero masses of neutrinos [20, 43].

Another significant limitation of the model is its treatment of the matter-antimatter asymmetry in the universe. Although the SM allows for CP violation via the CKM matrix, which is necessary to explain the imbalance, the degree of CP violation it predicts is insufficient to account for the observed dominance of matter over antimatter. Furthermore, the model does not resolve the strong CP problem [44], which is the unexplained conservation of CP symmetry in QCD, despite its known violation in weak interactions.

Additionally, the standard model does not explain why there are exactly three generations of fermions, nor does it address the specific organization and hierarchy among these generations.

A critical issue not addressed by the SM is the hierarchy problem [45]. This problem stems from the unexpectedly low mass of the Higgs boson compared to the Planck scale ($\Lambda = 10^{19}\,\text{GeV}$), where gravitational forces become significant. Quantum field theory suggests that scalar particles like the Higgs should receive large quantum corrections to their masses from loop diagrams involving all interacting particles. These corrections, which would be of $\mathcal{O}(\Lambda^2)$, should elevate the Higgs mass to the Planck scale. Yet, the Higgs boson's observed mass of about 125 GeV suggests that an unnatural fine-tuning is required to cancel out the corrections and keep the Higgs boson mass at the electroweak scale.

Many beyond-the-standard model (BSM) theories have thus been proposed to address some or all of the unresolved issues of the SM. The BSM theories relevant to the study presented in this thesis are discussed in detail in Sec. 2.3.

2.2 The Higgs Boson at the LHC

Following the discovery of the Higgs boson, the high-energy physics community has concentrated on thoroughly characterizing its properties. Efforts have been dedicated to ensuring its consistency with the standard model while also investigating potential new physics phenomena that the Higgs boson might reveal.

At the LHC, the SM Higgs boson can be produced primarily through the channels illustrated in Fig. 2.4. The *gluon-gluon fusion* (ggF), shown in the top left of Fig. 2.4, forms the dominant Higgs production mode. As gluons are massless, they cannot directly couple to the Higgs boson; instead, the interaction occurs via a quark loop initiated by the two initial state gluons. The dominant loop contribution arises from the top quark due to its large mass. This mechanism is particularly sensitive to BSM effects, as potential BSM particles could alter the quark loop, thus affecting the observed cross-section of this process. The second most prominent Higgs production mechanism is the *vector boson fusion* (VBF or qqH). In this mechanism, the Higgs boson is produced through the fusion of two W or Z bosons, which are radiated by initial state quarks. The Feynman diagram for this production mode is illustrated in Fig. 2.4 at the top right. *Vector-boson-associated production* (VH), or Higgs-strahlung, is another crucial production mode where the Higgs boson is produced alongside a W or Z boson. This process occurs when a virtual vector boson decays into its on-shell state while emitting a Higgs boson, as shown in the bottom left plot in Fig 2.4. In the VH production, the Higgs boson tends to be boosted, recoiling against the high-momentum leptons and jets

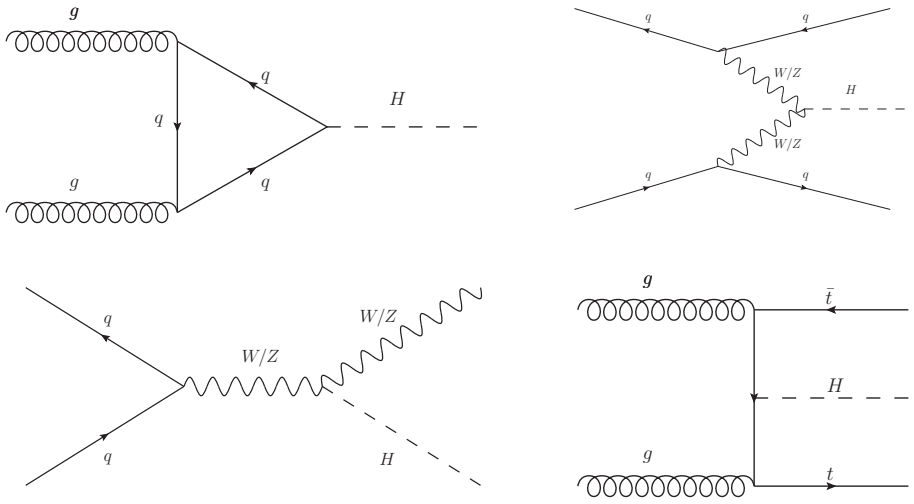

Figure 2.4: Feynman diagrams for the major production mechanisms of the Higgs at the LHC: gluon-gluon fusion (top left), vector boson fusion (top right), Higgs-strahlung (bottom left) and top-pair associated production (bottom right).

emanating from the vector bosons' decay. Finally, the Higgs boson can be produced via *top-quark pair associated production* ($t\bar{t}H$), as shown in Fig. 2.4 (bottom right). Though less dominant, this mode allows for the direct measurement of the top quark Yukawa coupling. A very similar Higgs production mechanism is the b-quark pair associated production $b\bar{b}H$. Though it has a comparable cross-section to $t\bar{t}H$, identifying the Higgs in $b\bar{b}H$ production is experimentally challenging due to the difficulty in distinguishing it from the background. Single-top associated production also contributes to Higgs production but is comparatively less dominant.

The SM Higgs production cross section predicted for the various mechanisms, for a Higgs mass of 125 GeV, is shown as a function of the center-of-mass energy, $\sqrt{s}$, of the proton-proton collisions in the left plot in Fig. 2.5.

The decay modes of the SM H include direct decays to massive particles and loop-induced decays to massless particles. At a Higgs boson mass, m_{H}, of 125 GeV, the most significant branching ratio is for the decay into a pair of b-quarks, followed by a pair of W bosons. The right plot in Fig. 2.5 illustrates the predicted branching ratios of the Higgs boson as a function of its mass.

Building on these theoretical predictions, extensive efforts have been undertaken by the CMS and ATLAS Collaborations to verify the consistency of experimental data with the SM [9,10]. A wide range of production mechanisms and decay modes have been extensively probed by the two experiments. Figure 2.6 showcases the results obtained by the CMS Collaboration by analyzing data collected at $\sqrt{s} = 13$ TeV. The two panels report the agreement of the experimental measurements with respect to the SM prediction in the form of signal strength modifiers μ. In this context, a perfect agreement with the SM corresponds to the value $\mu = 1$. The left panel displays the signal strength modifiers for six different production mechanisms

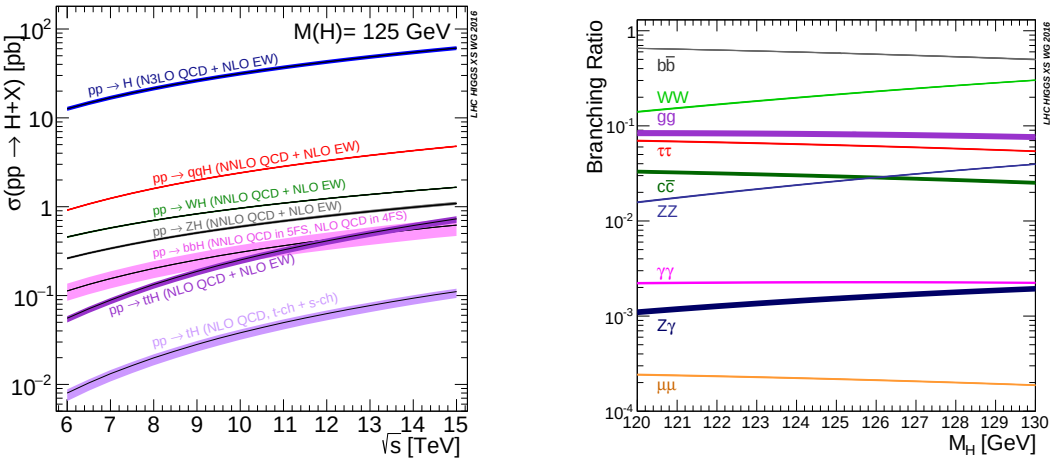

Figure 2.5: (Left) The SM Higgs production cross-sections for different mechanisms as a function of the center-of-mass energy of the proton-proton collisions, for a Higgs boson of mass 125 GeV. (Right) The SM Higgs boson branching ratios as a function of the Higgs boson [46].

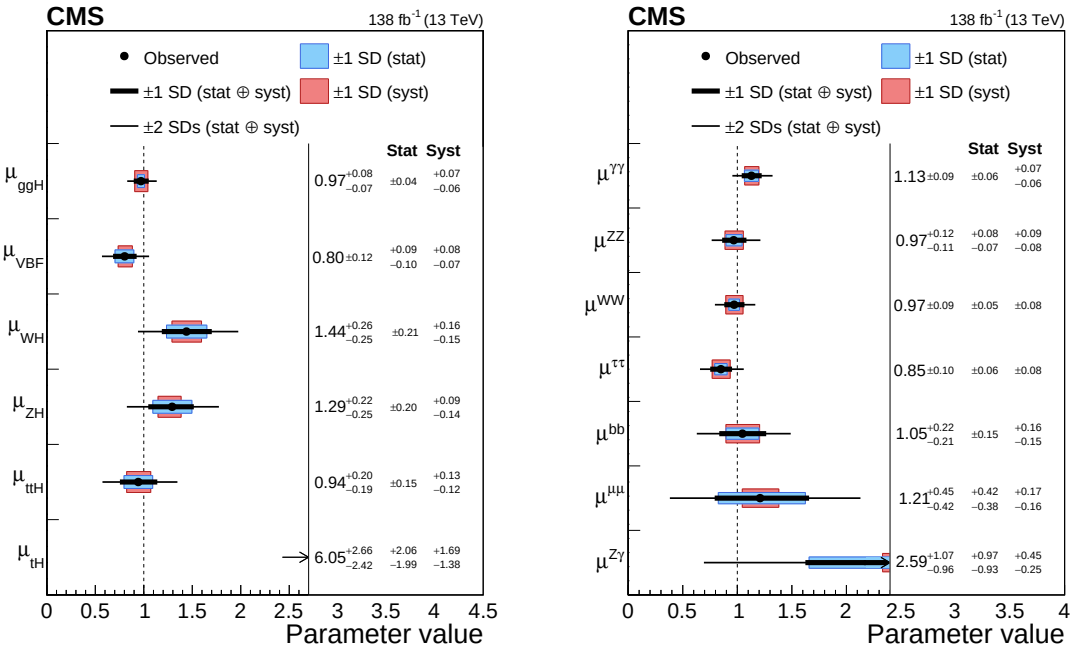

Figure 2.6: Signal strength parameters extracted for various production modes μ_i assuming $B_i = B_i^{\mathrm{SM}}$ (left), and decay channels μ_f assuming $\sigma_i = \sigma_i^{\mathrm{SM}}$ (right). The vertical dashed line at unity indicates perfect agreement with the SM) [9].

with decay branching ratios set to SM predictions; the right panel shows the signal strength for seven decay modes, assuming production cross sections are fixed at SM-predicted values. Most observed values of μ are close to unity or agree with it within uncertainties, indicating a strong overall agreement with the SM[3].

Despite this apparent agreement, there is still much to explore. Although experimental results largely match the predictions of the SM, significant gaps in the understanding of Higgs physics remain. These gaps could lead to discoveries of new phenomena beyond the SM. Notably, the decays of the Higgs boson to second-generation fermions, such as muons[4] and strange quarks, have not yet been observed. Any deviations in the branching ratios of these decays from SM predictions could signal modifications indicative of BSM physics. Additionally, further precision measurements of Higgs properties could uncover deviations in areas such as the couplings and mass, potentially pointing to new interactions or particles that BSM theories postulate. Another area of particular interest is the Higgs self-coupling; its measurement could reveal alterations in the Higgs potential, suggesting new dynamics at the electroweak scale that could influence our understanding of the universe's fundamental structure.

Furthermore, measurements of the Higgs boson couplings, as reported in Ref. [9,10], indicate that significant branching fractions for Higgs decays to beyond-SM particles remain a viable possibility, with ATLAS and CMS establishing upper limits at a 95% confidence level of 12% and 16%, respectively. These findings highlight the critical importance of continuing the search for "exotic" decays of the Higgs boson, which could potentially uncover new aspects of particle physics. Moreover, this possibility, combined with various other limitations of the standard model, serves as a compelling motivation to explore BSM theories with extended Higgs sectors.

2.3 Higgs Beyond the Standard Model

Many BSM theories aim to extend the SM framework in a way that retains its well-tested predictions while expanding its theoretical scope to address the critical shortcomings mentioned in Sec. 2.1.8. One popular approach within these theories involves the extension of the scalar sector beyond the single Higgs boson of the SM. Such extensions not only address critical gaps in the SM but also predict novel phenomena.

The presence of additional scalar fields modifies the Higgs potential, enabling new interactions within the Higgs sector. These modifications can facilitate the decay of the Higgs boson into lighter scalar or pseudoscalar particles that are otherwise not present in the standard model. Such exotic decay processes offer a direct avenue to probe new physics.

Examples of such models include the Minimal Supersymmetric Standard Model (MSSM), which incorporates supersymmetry to pair every SM particle with a corresponding superpartner [49, 50]. This effectively doubles the particle count. The enriched spectrum of this model helps resolve the hierarchy problem by allowing fermion and boson loops to contribute

[3]The tH production rate deviates from the SM predictions, but the significance of this variation is not sufficient to confirm a definitive anomaly [47].

[4]The CMS Collaboration has reported evidence for Higgs boson decays to muons, though it has not yet reached the significance required for discovery [48].

opposite-sign corrections to the Higgs boson mass, which are eventually canceled out. The MSSM also includes two Higgs doublets, which result in a complex array of five physical Higgs bosons. Extending this further, the Next-to-Minimal Supersymmetric Standard Model (NMSSM) adds an additional complex singlet, enriching the Higgs sector and solving the μ problem of the MSSM by naturally generating an effective μ-term through the vacuum expectation value of the singlet [51,52]. This addition results in seven physical Higgs states—three CP-even, two CP-odd, and two charged Higgs bosons. It also enables new decay pathways, such as the SM-like Higgs decaying into two lighter CP-odd Higgs bosons, if kinematically allowed.

Similarly, the scalar-singlet extension of the standard model introduces a single scalar field that is a singlet under SM gauge groups. This singlet might be stable and invisible, potentially contributing to dark matter if it does not decay. If the singlet mixes with the SM Higgs field, the model can yield two physical Higgs states, one of which could be identified as the 125 GeV Higgs [53]. This model then allows for exotic decays of the Higgs boson into a pair of the other scalar particle, if kinematically feasible [54].

The Two-Higgs-Doublet Model with an additional complex scalar singlet (2HDM+S) [11,55] stands out as particularly significant for this discussion. It builds on the Two-Higgs-Doublet Model (2HDM) [56,57], which extends the Higgs sector by introducing a second Higgs doublet. The 2HDM naturally appears in many BSM frameworks, including the MSSM. The 2HDM and 2HDM+S, by extension, lead to a richer Yukawa sector, allowing for new sources of CP violation [56,58]. They also support a strongly first-order electroweak phase transition, essential for explaining the universe's matter-antimatter asymmetry through electroweak baryogenesis [59]. Some realizations of 2HDM+S also offer dark matter candidates, further enhancing its theoretical appeal [60]. The addition of a singlet in 2HDM+S further expands the Higgs spectrum, resulting in seven Higgs states in total that significantly enrich the Higgs phenomenology. In particular, this model naturally accommodates exotic Higgs decays, where the Higgs boson decays into a pair of lighter scalars or pseudoscalars. The NMSSM can be viewed as a specific realization of 2HDM+S. However, the more general 2HDM+S framework provides greater flexibility in Higgs interactions and couplings, offering richer phenomenology and making it especially relevant for experimental searches targeting such decays.

Hence, the following subsection provides a detailed discussion of the Higgs sector in the 2HDM+S and the exotic decay possibilities it offers. Since the 2HDM forms the foundation of this model, the discussion begins with an introduction to the 2HDM before exploring the impact of the additional singlet and the resulting phenomenology.

2.3.1 The Two Higgs Doublet Model + Singlet

One of the simplest possible extensions to the standard model involves the addition of an extra doublet to its minimal Higgs sector. This approach gives rise to the general class of models known as the Two-Higgs-Doublet Models.

The most general 2HDM Higgs potential, assuming CP conservation, is given by:

$$\begin{aligned} V_{\text{2HDM}} = & m_{11}^2(\Phi_1^\dagger\Phi_1) + m_{22}^2(\Phi_2^\dagger\Phi_2) - m_{12}^2(\Phi_1^\dagger\Phi_2 + \Phi_2^\dagger\Phi_1) \\ & + \frac{\lambda_1}{2}(\Phi_1^\dagger\Phi_1)^2 + \frac{\lambda_2}{2}(\Phi_2^\dagger\Phi_2)^2 + \lambda_3(\Phi_1^\dagger\Phi_1\Phi_2^\dagger\Phi_2) \\ & + \lambda_4(\Phi_1^\dagger\Phi_2\Phi_2^\dagger\Phi_1) + \frac{\lambda_5}{2}((\Phi_1^\dagger\Phi_2)^2 + (\Phi_2^\dagger\Phi_1)^2) \end{aligned} \tag{2.44}$$

where m_{11}, m_{22}, m_{12}, λ_1, λ_2, λ_3, λ_4 and λ_5 are real parameters. Φ_1 and Φ_2 are the two Higgs doublets of the form:

$$\Phi_i = \begin{pmatrix} \phi_i^+ \\ \phi_i^0 \end{pmatrix}, \; i = 1, 2. \tag{2.45}$$

Following minimization, the two Higgs doublets $\Phi_{1,2}$ acquire vacuum expectation values $v_{1,2}$. These VEVs contribute jointly to the electroweak symmetry breaking, i.e., $\sqrt{v_1^2 + v_2^2} = v \approx 246\,\text{GeV}$, thus establishing a fundamental connection between the vacuum expectation values in the 2HDM and the SM. The parameter $\tan\beta$ is often introduced as the ratio of these two VEVs:

$$\tan\beta = \frac{v_2}{v_1}. \tag{2.46}$$

The 2HDM results in five physical Higgs states, which include:

- Two charged Higgs bosons $H^\pm$, which arise from the charged components of the doublets;
- Two CP-even neutral bosons, h and H^0, which result from the mixing of neutral components of Φ_1 and Φ_2, with α representing the mixing angle. By convention, h is taken to be lighter than H^0;
- One CP-odd neutral boson, A, which results from the imaginary parts of the neutral components of the doublets.

In 2HDM, flavor-changing neutral currents (FCNCs) are not forbidden at the tree level. However, they are strongly suppressed experimentally. Unlike the SM, which naturally suppresses FCNCs, the 2HDM requires additional mechanisms to control these unwanted currents. A commonly implemented solution is the Z_2 symmetry, which restricts each fermion group to couple exclusively to one of the Higgs doublets. This approach effectively allows for the definition of four distinct types of 2HDMs as summarized in Table. 2.1.

Type	up-type quarks	down-type quarks	charged leptons
Type I	Φ_2	Φ_2	Φ_2
Type II	Φ_2	Φ_1	Φ_1
Type III	Φ_2	Φ_2	Φ_1
Type IV	Φ_2	Φ_1	Φ_2

Table 2.1: Coupling of fermions to Higgs doublets in different types of 2HDM with a Z_2 symmetry.

	Coupling	Type I	Type II	Type III	Type IV
h	g_{hVV}	$\sin(\beta-\alpha)$	$\sin(\beta-\alpha)$	$\sin(\beta-\alpha)$	$\sin(\beta-\alpha)$
	g_{huu}	$\cos\alpha/\sin\beta$	$\cos\alpha/\sin\beta$	$\cos\alpha/\sin\beta$	$\cos\alpha/\sin\beta$
	g_{hdd}	$\cos\alpha/\sin\beta$	$-\sin\alpha/\cos\beta$	$\cos\alpha/\sin\beta$	$-\sin\alpha/\cos\beta$
	$g_{h\ell\ell}$	$\cos\alpha/\sin\beta$	$-\sin\alpha/\cos\beta$	$-\sin\alpha/\cos\beta$	$\cos\alpha/\sin\beta$
H^0	g_{H^0VV}	$\cos(\beta-\alpha)$	$\cos(\beta-\alpha)$	$\cos(\beta-\alpha)$	$\cos(\beta-\alpha)$
	g_{H^0uu}	$\sin\alpha/\sin\beta$	$\sin\alpha/\sin\beta$	$\sin\alpha/\sin\beta$	$\sin\alpha/\sin\beta$
	g_{H^0dd}	$\sin\alpha/\sin\beta$	$\cos\alpha/\cos\beta$	$\sin\alpha/\sin\beta$	$\cos\alpha/\cos\beta$
	$g_{H^0\ell\ell}$	$\sin\alpha/\sin\beta$	$\cos\alpha/\cos\beta$	$\cos\alpha/\cos\beta$	$\sin\alpha/\sin\beta$
A	g_{AVV}	0	0	0	0
	g_{Auu}	$\cot\beta$	$\cot\beta$	$\cot\beta$	$\cot\beta$
	g_{Add}	$-\cot\beta$	$\tan\beta$	$-\cot\beta$	$\tan\beta$
	$g_{A\ell\ell}$	$-\cot\beta$	$\tan\beta$	$\tan\beta$	$-\cot\beta$

Table 2.2: Couplings of the neutral scalar and pseudoscalar mass eigenstates in the four types of 2HDM with a Z_2 symmetry. The couplings are normalized to those of the SM Higgs. Here, the subscripts V, u, d, and ℓ correspond to the vector bosons, up-type quarks, down-type quarks, and leptons, respectively.

In the Type I 2HDM, all fermions exclusively couple to the second Higgs doublet, Φ_2. Conversely, in the Type II model, which resembles the MSSM, down-type quarks and leptons couple to Φ_1, while up-type quarks couple to Φ_2. The Type III model differentiates the couplings by assigning leptons to Φ_1 and quarks to Φ_2. Finally, in Type IV, down-type quarks couple to Φ_1, whereas both up-type quarks and leptons couple to Φ_2.

The couplings of the neutral Higgs bosons to the SM fermions and gauge fields, relative to the SM values, are summarized in Table 2.2. These couplings are governed by the angles α and β. The light CP-even scalar h interacts with gauge bosons with a relative strength of $\sin(\beta-\alpha)$, whereas the heavier scalar H^0 has couplings proportional to $\cos(\beta-\alpha)$. The CP-odd Higgs boson A does not couple to W and Z gauge bosons at tree level due to its pseudoscalar nature. The couplings of the neutral Higgs states to the gauge bosons are uniform across all 2HDM types. For fermions, however, the interaction patterns are significantly more complex, as they depend not only on the 2HDM type but also on the fermion type involved.

Theoretically, 2HDMs could allow for exotic decays of the 125 GeV Higgs boson, such as $h \to AA$, $H^0 \to hh, AA$ or $h \to \mathrm{Z}A$ with the daughter (pseudo)scalars subsequently decaying into SM fermions or gauge bosons. In this context, the 125 GeV state is identified with either h or H^0. However, current experimental data impose stringent constraints on the parameter space of 2HDMs, significantly limiting the possibilities for observing such exotic Higgs decay phenomenology within this model framework.

To address these constraints and expand the phenomenological possibilities, a strategic extension can be implemented within the 2HDM framework by introducing a complex scalar singlet. This modification leads to the formulation of the Two-Higgs-Doublet Model with a singlet (2HDM+S). The addition of the complex singlet introduces new dynamics and

interactions, including new exotic decay channels for the 125 GeV Higgs.

The complex singlet S is defined as:

$$S = \frac{1}{\sqrt{2}} S_R + i S_I, \tag{2.47}$$

where S_R and S_I are the real and imaginary parts of the singlet, respectively. The singlet introduces the following additional terms to the scalar potential [55]:

$$\begin{aligned} V_S = & (\xi S + h.c.) + m^2 S^\dagger S + \left(\frac{m'_S}{2} S^2 + h.c. \right) \\ & + \left(\frac{\mu_{S1}}{6} S^3 + h.c. \right) + \left(\frac{\mu_{S2}}{2} S S^\dagger S + h.c. \right) \\ & + \left(\frac{\lambda''_1}{24} S^4 + h.c. \right) + \left(\frac{\lambda''_2}{6} S^2 S^\dagger S + h.c. \right) + \left(\frac{\lambda''_3}{4} (S^\dagger S)^2 \right) \\ & + \left[S(\mu_{11} \Phi_1^\dagger \Phi_1 + \mu_{22} \Phi_2^\dagger \Phi_2 + \mu_{12} \Phi_1^\dagger \Phi_2 + \mu_{21} \Phi_2^\dagger \Phi_1) + h.c. \right] \\ & + S^\dagger S \left[\lambda'_1 \Phi_1^\dagger \Phi_1 + \lambda'_2 \Phi_2^\dagger \Phi_2 + (\lambda'_3 \Phi_1^\dagger \Phi_2 + h.c.) \right] \\ & + S^2 \left[\lambda''_4 \Phi_1^\dagger \Phi_1 + \lambda''_5 \Phi_2^\dagger \Phi_2 + \lambda''_6 \Phi_1^\dagger \Phi_2 + \lambda''_7 \Phi_2^\dagger \Phi_1 + h.c. \right], \end{aligned} \tag{2.48}$$

such that the full general scalar potential for a 2HDM+S model is given by:

$$V_{\text{2HDM+S}} = V_{\text{2HDM}} + V_S. \tag{2.49}$$

In this model, the electroweak symmetry breaking condition $\sqrt{v_1^2 + v_2^2} = v \approx 246\,\text{GeV}$ is maintained, ensuring that the Higgs doublets Φ_1 and Φ_2 correctly contribute to the masses of the W and Z bosons. The vacuum expectation value of the singlet, v_S, does not affect these gauge boson masses as it does not interact with the electroweak gauge fields.

The addition of a complex singlet scalar enhances the Higgs sector in the 2HDM+S model compared to the standard 2HDM, increasing the total count of physical Higgs states to seven. The real part of the singlet S mixes with the neutral CP-even components of the doublets to yield three CP-even Higgs states denoted h_1, h_2, and h_3. With minimal mixing between the singlet and the doublet components, two of the CP-even Higgs states, usually h_1 and h_2, retain properties similar to the h and H^0 bosons of the 2HDM. The additional state, h_3, primarily emerges from the singlet's real component. Furthermore, the imaginary part of the singlet can mix with the CP-odd component from the doublets, yielding two CP-odd mass eigenstates. Here, the new pseudoscalar a_1 is mainly influenced by the singlet's imaginary part while a_2 retains the properties of the standard 2HDM's CP-odd Higgs, A. The addition of the singlet does not affect the charged Higgs bosons, maintaining their count at two ($H^\pm$) from the doublets.

By regarding the 2HDM component of the 2HDM+S to be in the decoupling limit, the lightest CP-even Higgs boson, h, can be identified with the observed 125 GeV Higgs, H. The very

small mixing considered in the 2HDM+S ensures that $h_1 \equiv h$, and the SM-like nature of h is not spoiled. In this framework, the 2HDM+S introduces possible exotic decays of the 125 GeV Higgs in the form of $H \to h_3 h_3$ and $H \to a_1 a_1$, subject to kinematic feasibility.

This thesis primarily considers the decay channel $H \to a_1 a_1$, as detailed in Chapter 5. This channel is favored over $H \to h_3 h_3$ due to its simpler coupling structure, which allows for more straightforward interpretations. The pseudoscalar a_1 mixes with only one component from the doublets, unlike the scalar h_3, which mixes with two components. This single mixing parameter in the pseudoscalar case reduces the complexity compared to the scalar case, which involves an additional degree of freedom from the extra mixing parameter.

The mostly-singlet-like pseudoscalar a_1 can be represented as:

$$a_1 = \cos\theta S_I + \sin\theta A, \;\; \theta \ll 1, \tag{2.50}$$

where θ is the mixing angle that controls the composition of the pseudoscalar a_1. By tuning the parameters of the model, a_1 can be made lighter than the 125 GeV Higgs boson, kinematically allowing decays of the form $H \to a_1 a_1$ for $m_{a_1} < \frac{m_h}{2} \approx 62.5$ GeV.

Unlike the Higgs doublets Φ_1 and Φ_2, the singlet S does not couple directly to the standard model fermions. Instead, it interacts with the fermion sector indirectly through its mixing with the Higgs doublets. Consequently, the pseudoscalar a_1 decays into SM fermions via the existing couplings of A (detailed in Table 2.2) scaled by $\sin\theta$. This facilitates exotic decays of the form $H \to a_1 a_1 \to X\overline{X}Y\overline{Y}$, where X and Y are standard model fermions.

The 2HDM+S model retains the four canonical types of fermion couplings characteristic of the 2HDM, each determining how the pseudoscalars interact with SM fermions. Consequently, the specific branching ratios of a_1 decaying into SM particles are highly dependent on the selected type of 2HDM+S. Moreover, the value of $\tan\beta$ and the mass of the pseudoscalar also significantly impact these decays. The different decay patterns showcased by each type of 2HDM+S are highlighted below:

- Type I: As all fermions couple to Φ_2, the branching ratios are independent of $\tan\beta$. Hence, the pseudoscalar couplings are equivalent to those in the SM. In this model, decays to heavier fermions are always favored whenever kinematically allowed.
- Type II: Here, down-type quarks and leptons couple to Φ_1, while up-type quarks couple to Φ_2. Consequently, for $\tan\beta > 1$, the decays to down-type quarks and leptons are enhanced, making $a_1 \to \tau\tau$ decays dominant above the tau-pair production threshold up to a pseudoscalar mass of ~ 10 GeV. Beyond this mass, decays to $b\overline{b}$ become accessible and significant. Conversely, for $\tan\beta < 1$, despite their lighter mass compared to tau or bottom quarks, decays to charm quarks become the dominant decay channel for the pseudoscalar a_1. The Type II 2HDM+S essentially represents the scalar sector of NMSSM.
- Type III: In this model, both up- and down-quarks couple to Φ_1 and leptons couple to Φ_2. The branching ratios to leptons (up-type and down-type quarks) are enhanced (suppressed) when $\tan\beta > 1$ across all pseudoscalar masses (in the kinematically allowed region). This channel is particularly favorable for $a_1 \to \tau\tau$ decays. The decays to quarks are only accessible for $\tan\beta < 1$.

- Type IV: Here, up-type quarks and leptons couple to Φ_1, while down-type quarks couple to Φ_2. The branching ratio to down-type quarks is enlarged for $\tan\beta > 1$, allowing decays of the type $a_1 \to b\bar{b}$, whereas decays $a_1 \to c\bar{c}/\tau\tau$ are only prominent for $\tan\beta < 1$.

Figure 2.7 illustrates the branching fractions for $a_1 \to X\overline{X}$ decays, where X represents SM particles, across different 2HDM+S models for specified $\tan\beta$ values.

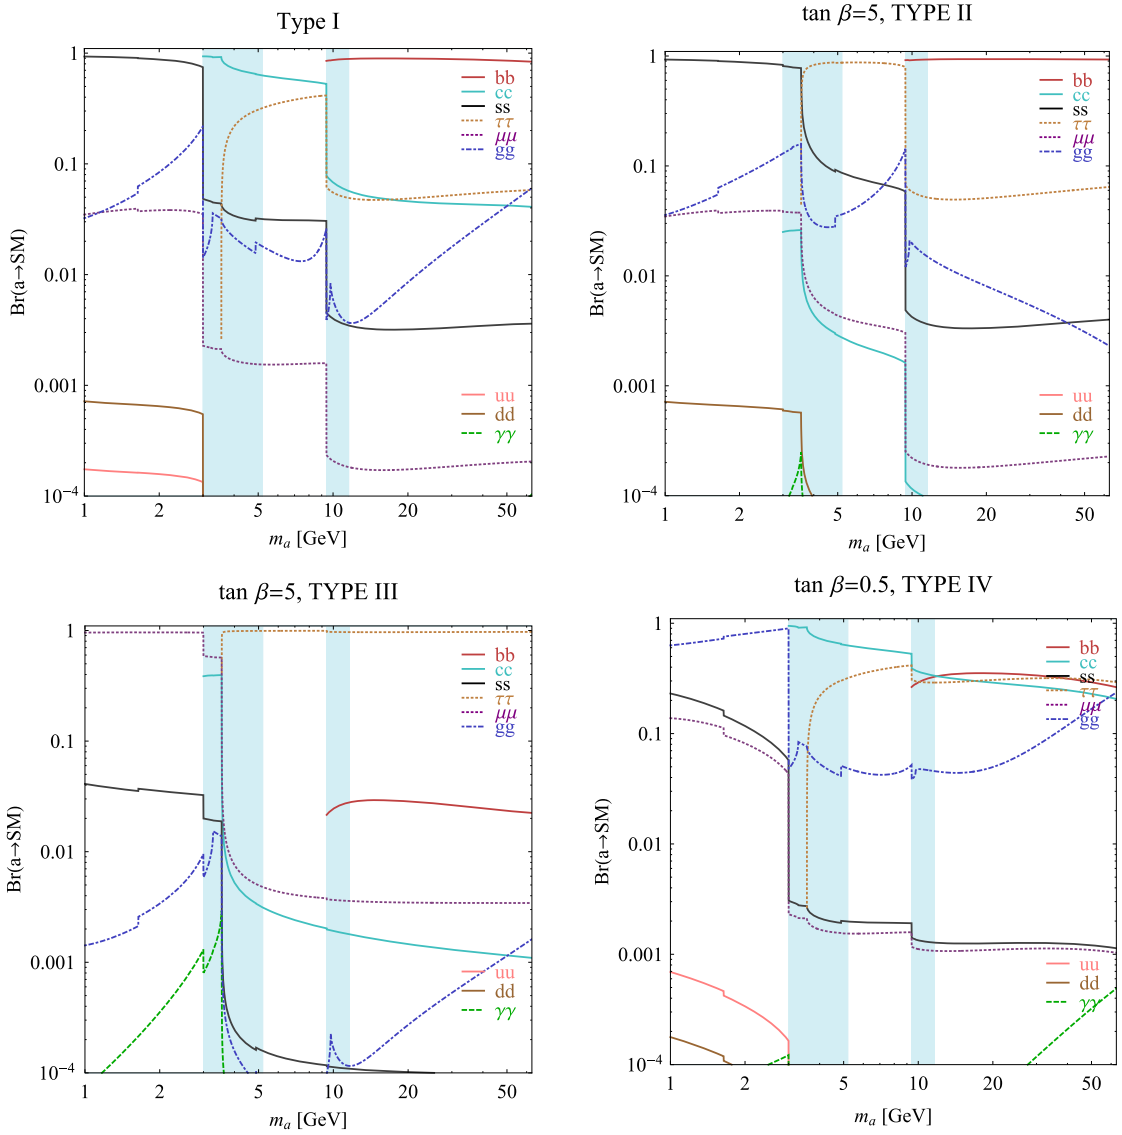

Figure 2.7: Branching ratios of the singlet-like pseudoscalar to a pair of SM particles for different 2HDM+S scenarios: Type I (top left), Type II $\tan\beta = 5$ (top right), Type III $\tan\beta = 5$ (bottom left) and Type IV [11].

As explained above, the leptonic decay $a_1 \to \tau\tau$ is highly favored in a considerable portion of the phase space of 2HDM+S. Depending on the specific model type, decays to tau leptons can even surpass those to b-quarks for pseudoscalar masses above the b-quark pair threshold. This underscores why, despite the challenging decay topology, it remains crucial to investigate pseudoscalars in final states involving tau leptons.

2.3.2 Current Constraints on 2HDM+S from Exotic Higgs Decay Searches

The Two-Higgs-Doublet Model plus singlet is of particular interest at the LHC as a well-motivated framework for exploring exotic decays of the 125 GeV Higgs boson. Experiments like ATLAS and CMS have extensively explored various decay channels for the $H \to a_1 a_1$ process, targeting final states such as fermion pairs (e.g., $b\bar{b}$, $\tau^+\tau^-$) and diphotons. However, no concrete evidence for these additional Higgs bosons or exotic Higgs decays has been observed thus far. These results have imposed stringent constraints on the model's parameter space, particularly on the cross section times branching fraction $\sigma\mathcal{B}(H \to a_1 a_1)$, relative to the standard model Higgs production cross section σ_{SM}, as a function of $\tan\beta$, and the pseudoscalar mass. Figures 2.8–2.10 illustrate the current exclusion limits for selected benchmark scenarios set by the CMS Collaboration, while constraints set by the ATLAS Collaboration can be viewed at [61].

Ongoing studies and future LHC runs, benefiting from increased luminosity and improved detection methods, are expected to further refine these constraints. This thesis contributes to these efforts by specifically analyzing the $H \to a_1 a_1 \to 4\tau$ channel and setting improved exclusion limits within the 2HDM+S framework.

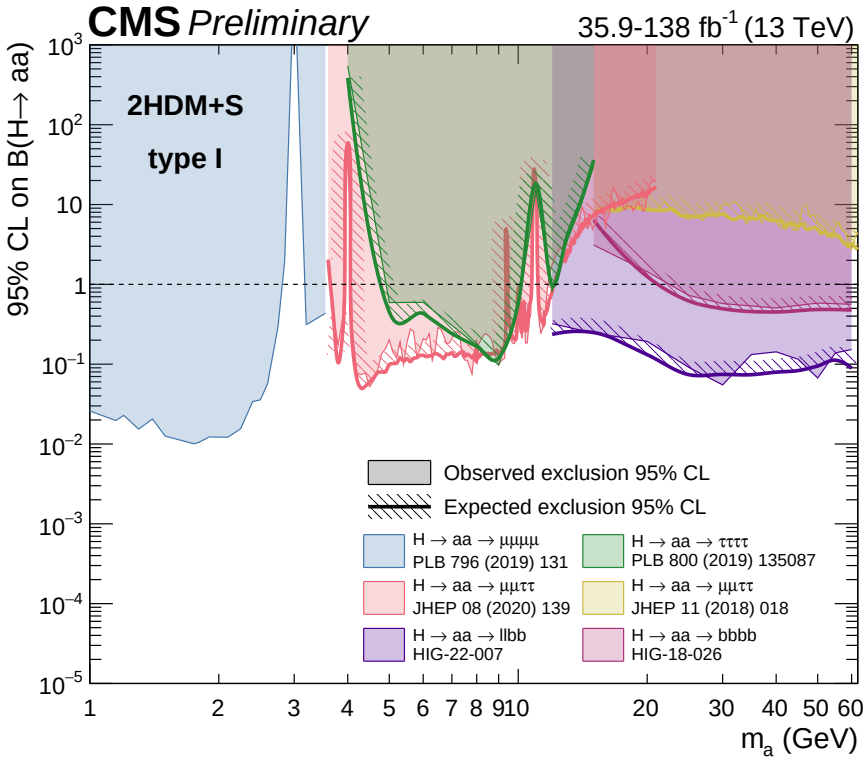

Figure 2.8: Observed and expected 95% confidence level upper limits on $\frac{\sigma}{\sigma_{SM}}\mathcal{B}(H \to a_1 a_1)$ in the 2HDM+S type-I scenario, set by the CMS experiment [62]. *Note that this thesis provides improved constraints for the* $H \to a_1 a_1 \to 4\tau$ *channel, which are not yet updated in the current figure.*

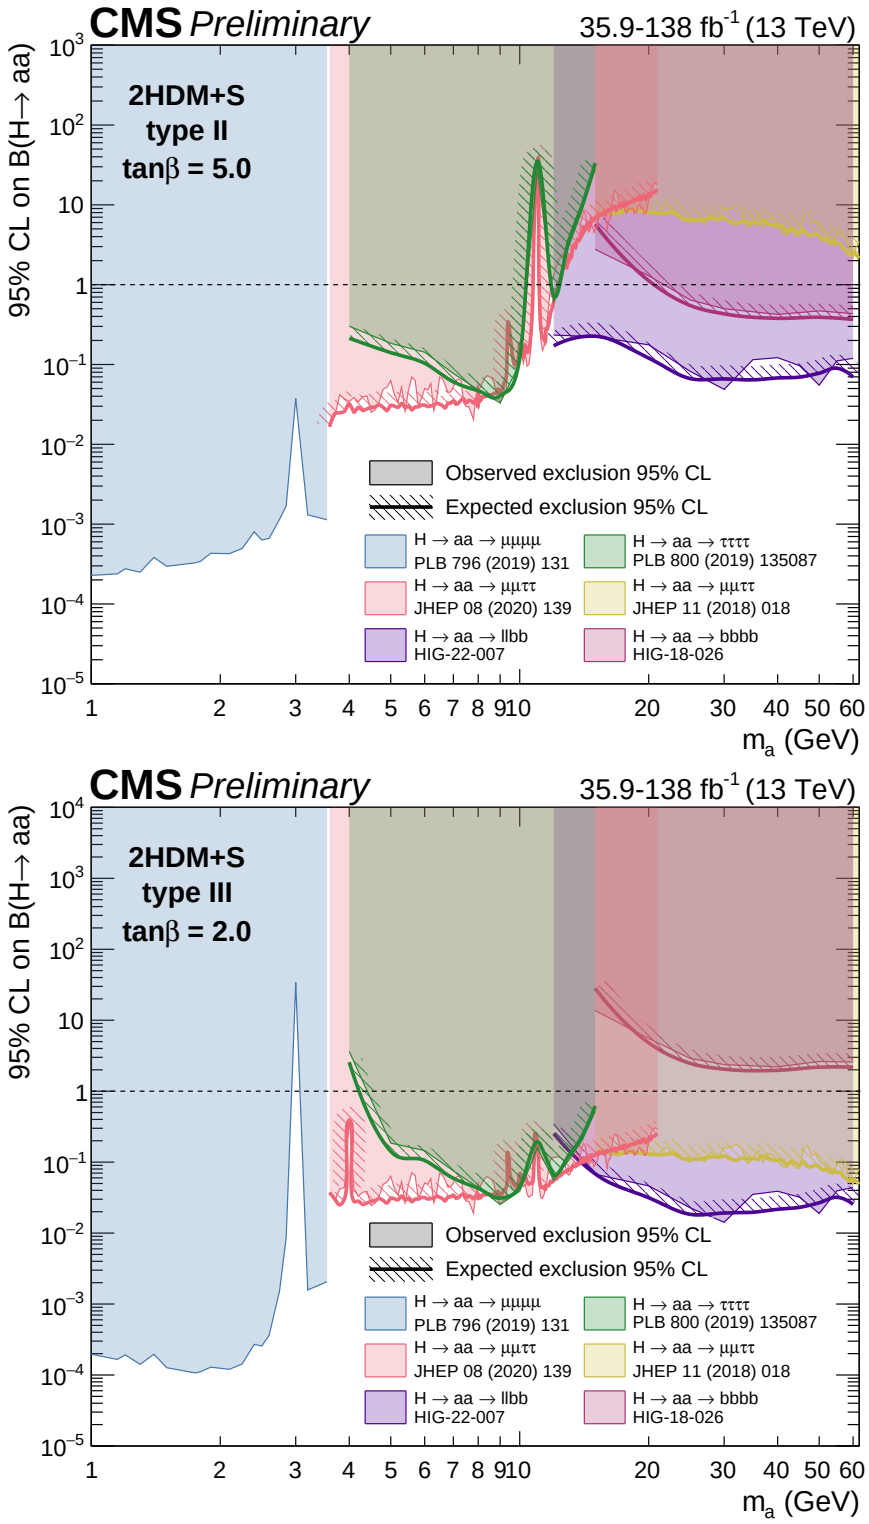

Figure 2.9: Observed and expected 95% confidence level upper limits on $\frac{\sigma}{\sigma_{\mathrm{SM}}}\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$ in the 2HDM+S type-II $\tan\beta$=5 scenario (top) and 2HDM+S type-III $\tan\beta$=2 scenario (bottom), set by the CMS experiment [62]. *Note that this thesis provides improved constraints for the* $\mathrm{H} \to \mathrm{a_1 a_1} \to 4\tau$ *channel, which are not yet updated in the current figure.*

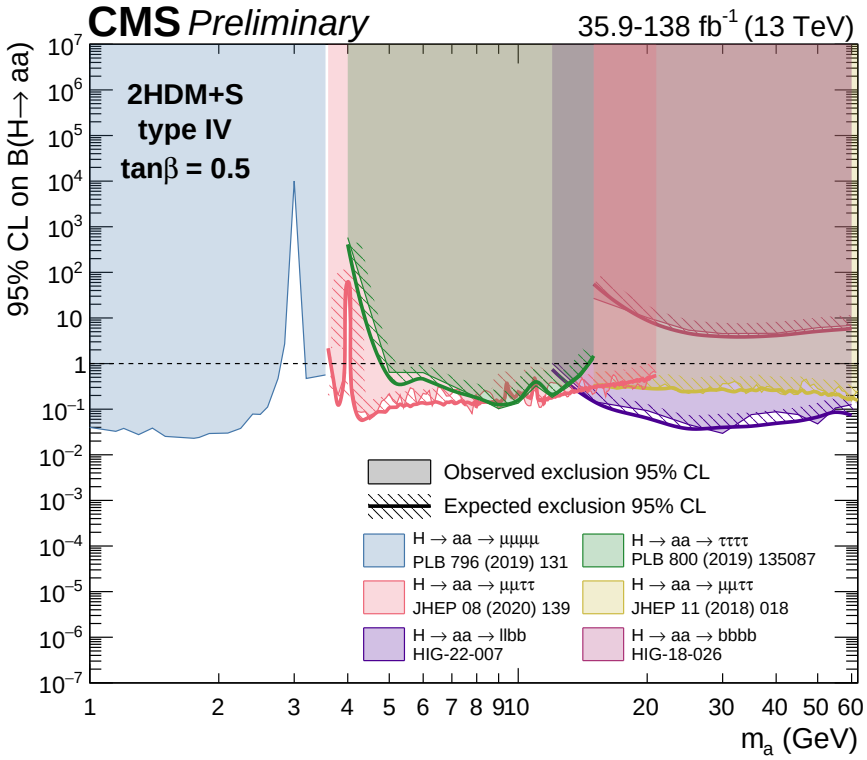

Figure 2.10: Observed and expected 95% confidence level upper limits on $\frac{\sigma}{\sigma_{\mathrm{SM}}}\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$ in the 2HDM+S type-IV $\tan\beta$=0.5 scenario, set by the CMS experiment [62]. *Note that this thesis provides improved constraints for the* $\mathrm{H} \to \mathrm{a_1 a_1} \to 4\tau$ *channel, which are not yet updated in the current figure.*

Chapter 3

Experimental Facilities

Contents

3.1 The Large Hadron Collider 33
3.1.1 The Accelerator Complex 34
3.1.2 Experiments at the LHC 35
3.1.3 LHC Design and Specifications 36
3.1.4 Operation Schedule 38
3.2 The Compact Muon Solenoid Experiment 41
3.2.1 Coordinate system 41
3.2.2 Detector components 42
3.2.3 Phase 2 upgrade of the CMS detector 57

The data analyzed in the context of this thesis has been collected by the Compact Muon Solenoid (CMS) detector at the Large Hadron Collider (LHC). This chapter offers a comprehensive overview of the design and workings of both the LHC and CMS.

3.1 The Large Hadron Collider

The Large Hadron Collider, constructed and operated by the European Organization for Nuclear Research (CERN), is the largest and most powerful particle accelerator in the world at present [63]. It is installed in the same underground circular tunnel that previously housed the Large Electron-Positron Collider (LEP) [64]. Having a total circumference of 26.7 km, this tunnel is located near Geneva, across the Franco-Swiss border, 50-170 m below the ground. The LHC is designed to collide protons at a maximum center-of-mass energy, $\sqrt{s} = 14$ TeV. It can also accelerate heavy ions, like lead and xenon, to a maximum energy of 2.76 TeV per nucleon. This section describes the design, functioning, and operational capabilities of the LHC, as well as the major experiments around it. As heavy ion collisions are outside the scope of this thesis, further discussions will be focused on proton-proton (pp) collisions.

3.1.1 The Accelerator Complex

The process of accelerating protons to very high energies involves multiple successive steps, with the LHC serving as the final stage in a complex chain of accelerators. The CERN accelerator complex is shown as a schematic in Fig. 3.1. The very first step in the acceleration journey is the ionization of hydrogen gas. Electrons are stripped away from hydrogen atoms by applying a strong electric field, leaving behind protons that are then bunched together. Till 2018, these proton bunches were injected into the Linear Accelerator 2 (Linac2) to carry out the first stage of acceleration [65]. The Linac2 is a 36 m long linear accelerator that uses radiofrequency (RF) cavities to accelerate protons up to an energy of 50 MeV. Although Linac2 was decommissioned in 2018, its successor, Linac4, has continued this role since 2020, delivering protons with energies around 160 MeV [66]. Once the protons have been accelerated by Linac2 (or Linac4), they are injected into the Proton Synchrotron Booster (PSB), a circular accelerator with a circumference of 160 m [67]. The PSB raises the energies of the protons to nearly 1.4 GeV. Following this, the protons enter a larger, 600 m long, circular accelerator called the Proton Synchrotron (PS) [68], where they are accelerated up to 25 GeV. The next in the chain is the Super Proton Synchrotron (SPS) [69], with a circumference of 7 km, which accelerates the protons to an energy of 450 GeV before injecting them into the LHC.

The single proton beam accelerated through the pre-accelerators is split into two at the SPS stage using magnetic separators and kicker magnets and injected into one of two antiparallel beam pipes at the LHC, where the protons then circulate in opposite directions under an ultra-high vacuum environment. The journey of the protons within the LHC is facilitated by a sophisticated system of superconducting electromagnets and RF cavities. Operating at 400 MHz, the 16 RF cavities in the LHC ensure progressive acceleration of the proton beams [71]. The LHC employs 1232 dipole magnets, each 15 m long, weighing 35 tonnes and generating magnetic fields of 8.3 T, to bend the beams along the circular path of the accelerator and keep them perfectly aligned with the orbit [72]. Additionally, quadrupole magnets, which have four magnetic poles arranged symmetrically around the beam pipe, are used to squeeze the beam either vertically or horizontally in order to focus it. There are 392 such quadrupole magnets in place, with each measuring 5 to 7 m in length [72]. Magnets of higher multipole orders are used to correct minor imperfections in the field geometry. The superconducting magnets are made from copper-stabilized niobium-titanium (NbTi) and are cooled to an extremely low temperature of 1.9 K (-271.25° C) through one of the largest and most sophisticated cryogenic systems in the world, utilizing superfluid helium [73].

Once the beams reach the desired energy and stability, they are further focused using special magnets and then collided at the interaction points (IPs). Three closely-spaced quadrupole magnets, creating a system called inner triplet, dramatically narrow the beams from 0.2 mm to just 16 μm at these points, thereby significantly increasing the chances of collisions. The LHC has four main interaction points along its circumference, namely IP1, IP2, IP5, and IP8. Each IP is equipped with an advanced particle detector, as explained further in Sec. 3.1.2.

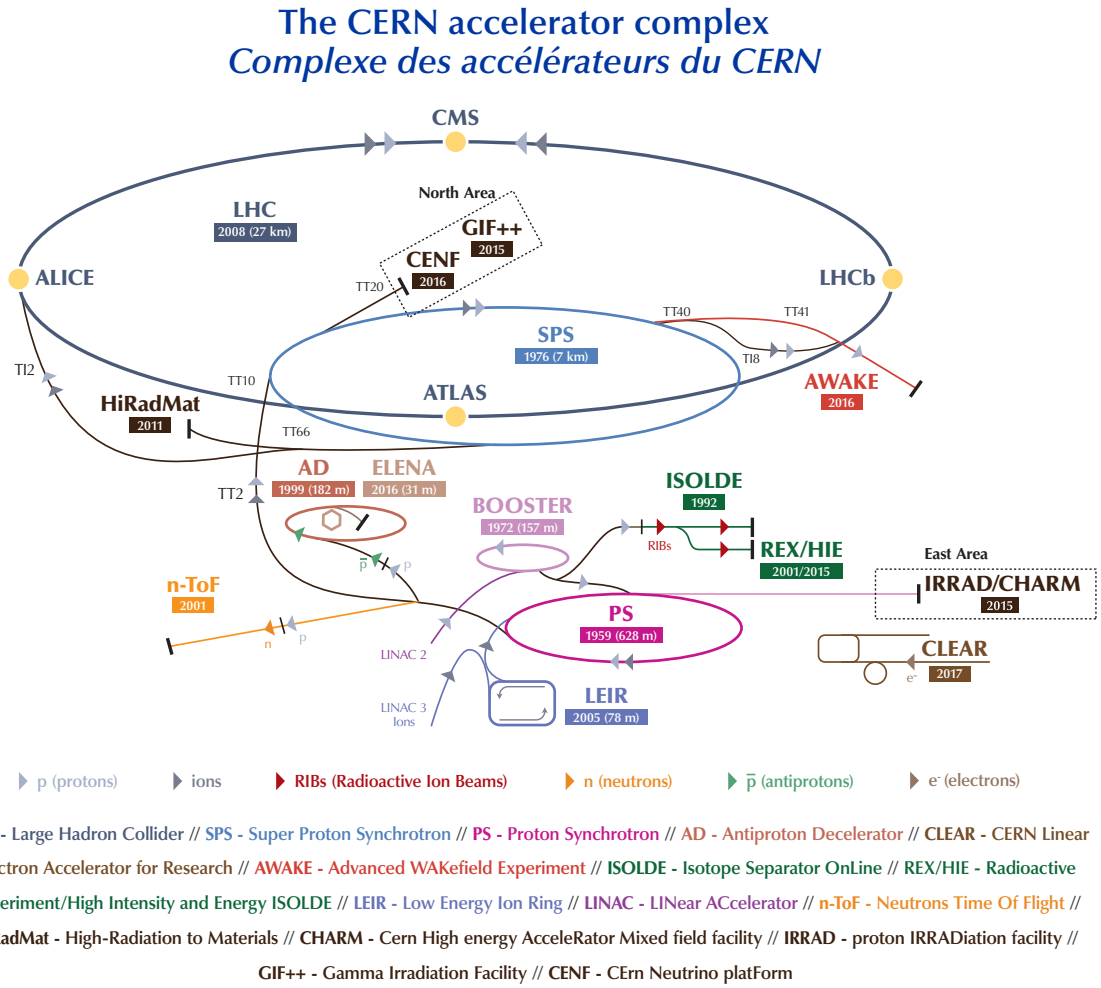

Figure 3.1: A schematic illustration of the accelerator complex at CERN as of 2018 [70]. The LHC is the last ring (dark blue) in a series of particle accelerators. The smaller machines are used in a specific order to accelerate proton beams. Since 2020, Linac4 has taken over the initial acceleration stage from Linac2.

3.1.2 Experiments at the LHC

Four major experiments are located along the LHC, with their detectors placed at the various interaction points. These experiments include:

1. **ATLAS** (A Toroidal LHC ApparatuS)
   Situated at IP1, ATLAS is one of the two general-purpose detectors at the LHC that explores a broad range of physics phenomena [74]. This 7000-tonne detector is the largest volume detector to be built to date.

2. **ALICE** (A Large Ion Collider Experiment)
   ALICE, located at IP2, is focused on heavy-ion physics and is specially designed to study quark-gluon plasma, a state of matter believed to have existed just after the Big Bang [75]. The experiment aims to provide insights into the dynamics of the strong force under

extreme energy densities and temperatures.

3. **CMS** (Compact Muon Solenoid)
   The CMS detector, housed at IP5, is the other general-purpose detector at the LHC [76]. Its design, which includes a compact but highly effective solenoid magnet system, enables precise measurements of particle properties.

4. **LHCb** (Large Hadron Collider beauty)
   Located at IP8, LHCb specializes in b-physics [77]. Its design features a single-arm forward spectrometer that effectively captures the decay products of B-hadrons, facilitating detailed studies of matter-antimatter asymmetries.

The position of these experiments with respect to the accelerator complex can be seen in Fig. 3.1.

Alongside these primary detectors, five other smaller experiments are located along the LHC ring. The Large Hadron Collider forward (LHCf) consists of two independent detectors located 140 m on either side of IP1 [78]. It aims to study particles that are scattered in the forward direction close to the beamline, which simulate cosmic ray interactions in the atmosphere. The Forward Search Experiment (FASER), operational since July 2022 and located 480 m downstream from the ATLAS IP along the beamline, searches for new light and weakly interacting particles that could shed light on dark matter [79]. The Scattering and Neutrino Detector at the LHC (SND@LHC), which began operation in mid-2022, is positioned 480 m from the ATLAS IP, slightly off the beamline; it detects collider neutrinos produced in the LHC collisions at small angles with respect to the beamline [80]. Additionally, the TOTal cross-section, Elastic scattering and diffraction dissociation Measurement (TOTEM) experiment at IP5 near CMS studies the properties of protons as they emerge from collisions at small angles [81]. The TOTEM experiment provides critical data to understand proton structure at high energies. Lastly, the Monopole & Exotics Detector At the LHC (MoEDAL), located at IP8 alongside LHCb, searches for magnetic monopoles and other highly ionizing particles, probing new physics beyond the SM [82].

3.1.3 LHC Design and Specifications

One of the main goals of the LHC is to observe and investigate rare processes that have much smaller cross sections compared to the total proton-proton inelastic cross section. These rare events occur infrequently, thereby requiring a high number of collisions to detect them. Luminosity, defined as the collision rate per unit area, has a crucial role in this context. A higher luminosity increases the number of collisions, thereby enhancing the probability of recording these rare events and improving statistical significance. The machine luminosity $\mathcal{L}$, also referred to as the instantaneous luminosity, depends solely on the beam parameters and can be mathematically expressed as:

$$\mathcal{L} = \frac{N_b^2 n_b f_{rev} \gamma_r}{4\pi \epsilon_n \beta^*} F, \tag{3.1}$$

where N_b is the number of particles per bunch, n_b the total number of bunches per beam, f_{rev} the beam revolution frequency, γ_r the relativistic gamma factor, ϵ_n the normalized transverse

Beam parameter	Description	Nominal value
N_b	Number of particles per bunch	1.15×10^{11}
n_b	Number of bunches	2808
f_{rev}	Revolution frequency	11.25 kHz
γ_r	Relativistic gamma	7461
ϵ_n	Normalised transverse emittance	3.75×10^{-6} mrad
β^*	Beta-value at IP	0.55 m
θ_c	Crossing angle at IP	285 μrad
σ_z	RMS bunch length	7.55 cm
σ^*	Transverse RMS beam size at IP	16.7 μm

Table 3.1: Nominal values of LHC beam parameters relevant for proton-proton collisions.

beam emittance, and β^* the beta function at the collision point. The geometric luminosity reduction factor, F, accounts for the effects of the crossing angle at the IP and is defined as:

$$F = \left(1 + \left(\frac{\theta_c \sigma_z}{2\sigma^*}\right)^2\right)^{-1/2}. \tag{3.2}$$

Here θ_c is the full crossing angle at the IP, σ_z the bunch length, and σ^* the transverse beam size at the IP [63]. The design values of the beam parameters are summarized in Table 3.1. To maximize the number of interactions per bunch crossing, β^* is reduced when approaching the IP. Additionally, θ_c is minimized in order to have proton bunches meet almost head-on.

The number of events generated per unit time for a particular physics process with cross section σ can be then estimated from the instantaneous luminosity as:

$$\frac{\partial N}{\partial t} = \mathcal{L} \cdot \sigma. \tag{3.3}$$

Integrated luminosity (L) gives a measure of the total number of collisions produced over a period of time and is defined as:

$$L = \int \mathcal{L} \partial t. \tag{3.4}$$

Increasing the instantaneous luminosity is essential for maximizing data collection. However, this increase also leads to an effect known as pileup (PU), wherein multiple proton-proton collisions occur within the same bunch crossing. The average number of pileup events per bunch crossing denoted $\langle \mathrm{PU} \rangle$, can be calculated based on the machine's luminosity $\mathcal{L}$ and the total inelastic cross-section of proton-proton interactions $\sigma_{\mathrm{in}}^{\mathrm{pp}}$, as:

$$\langle \mathrm{PU} \rangle = \frac{\mathcal{L} \cdot \sigma_{\mathrm{in}}^{\mathrm{pp}}}{n_b \cdot f_{rev}}. \tag{3.5}$$

The LHC was originally designed to achieve a peak instantaneous luminosity of $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$

and produce, on average, 25 pileup interactions. Details on the actual luminosity and pileup values reached throughout its operational period are discussed in Sec. 3.1.4.

3.1.4 Operation Schedule

The LHC operates in runs, which are periods of data taking typically spanning several years, followed by long shutdown (LS) periods for maintenance and upgrades. The operational schedule of the LHC is meticulously planned to maximize data collection while providing ample time for enhancements. Figure 3.2 gives a timeline of the LHC operations along with the milestones achieved (or planned) in terms of integrated and instantaneous luminosity [83]. The integrated and instantaneous luminosity delivered by the LHC to CMS during stable beams for pp collisions at nominal center-of-mass energies across the years is shown in Fig. 3.3.

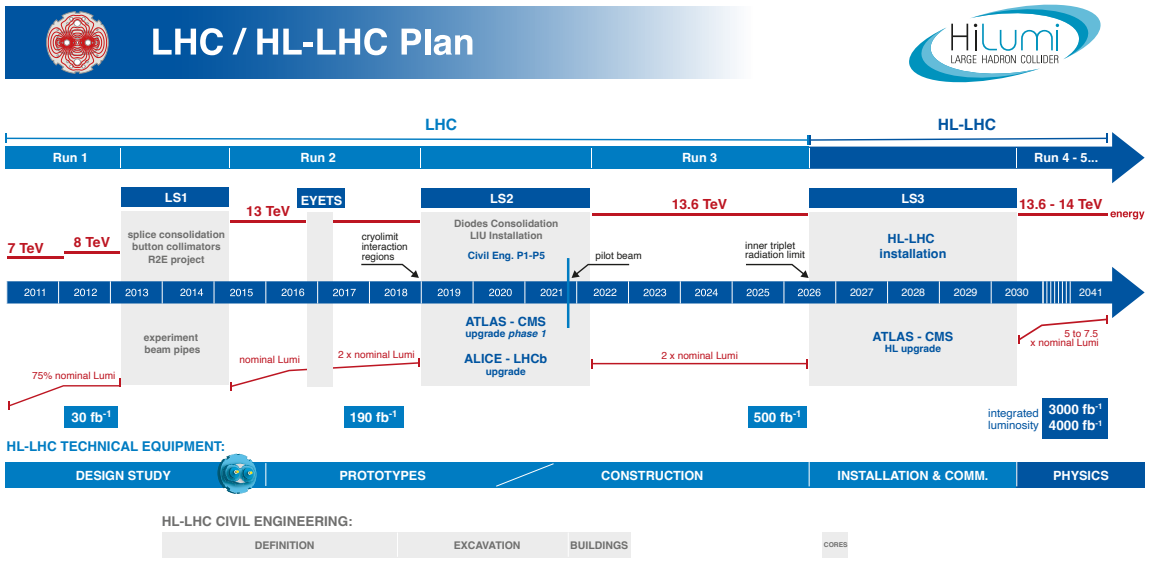

Figure 3.2: Schedule of the LHC operations. Also indicated are the maximum instantaneous and integrated luminosities achieved or planned for each run [83].

The first proton-proton collision at a center of mass energy of 7 TeV was recorded at the LHC on 30$^{\text{th}}$ March 2010, marking the beginning of a significant phase in high energy physics and the start of Run 1. The year 2010 was devoted to commissioning and establishing confidence in operational procedures. The official data taking for physics during Run 1 took place in 2011 and 2012. After progressively increasing the luminosity throughout 2011, the center of mass energy at the LHC was raised to 8 TeV on 5$^{\text{th}}$ April 2012. Run 1 accounts for 6.1 fb^{-1} of data collected at $\sqrt{s} = 7\,\text{TeV}$ and 23.3 fb^{-1} recorded at $\sqrt{s} = 8\,\text{TeV}$. It was pivotal in achieving the historic milestone of the discovery of the Higgs boson in 2012 by the ATLAS and CMS experiments independently [6–8]. After the completion of Run 1, the LHC entered its First Long Shutdown (LS1), lasting two years, wherein several efforts were undertaken to upgrade the LHC and the detectors to accommodate the increase of the energy per beam to 6.5 TeV and the heightened luminosity requirements planned for Run 2.

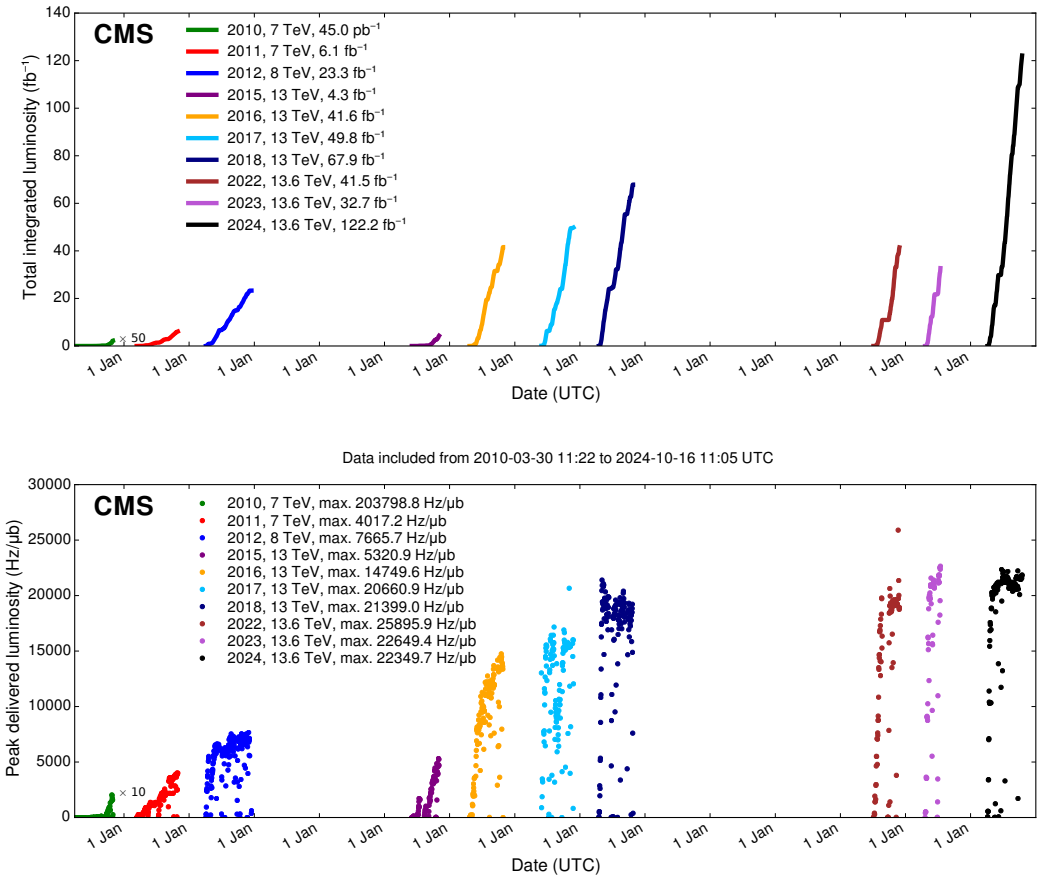

Figure 3.3: Integrated luminosity (top) and peak instantaneous luminosity (bottom) versus time, delivered to the CMS experiment during pp collisions for each year of data taking [84].

The LHC resumed operations in June 2015, with the start of Run 2. This marked yet another milestone for the LHC as pp collisions reached a record center of mass energy of 13 TeV. Run 2 data taking lasted from 2015 to 2018, with the operations in 2015 dedicated towards the commissioning of the new LHC configurations. The proton bunches were spaced apart by 25 ns, and a maximum peak instantaneous luminosity of $2.1\times10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, almost double the design value, was reached in Run 2. A total of around $160\,\mathrm{fb}^{-1}$ of data was delivered during this period, with which several crucial results were achieved, such as the first observation of the direct coupling of the Higgs boson to τ leptons [85] and b-quarks [86] as well as the first evidence of the Higgs boson coupling to muons [48]. After the conclusion of Run 2 in December 2018, the LHC's three-year-long Long Shutdown 2 (LS2) commenced, during which renovation works were carried out to further push the performance of the LHC and the detectors.

Run 3 of the LHC started in 2022 and is set to conclude in mid-2026. A record high center of mass energy was yet again reached with $\sqrt{s} = 13.6\,\mathrm{TeV}$. At the time of writing this thesis, the pp collision data taking of 2024 has concluded. Almost $196\,\mathrm{fb}^{-1}$ of data has been delivered by the LHC so far during Run 3, with the peak instantaneous luminosity slightly

surpassing that of Run 2. By the end of Run 3, around $500\,\text{fb}^{-1}$ is expected to be delivered to the experiments, providing a large dataset to better probe rare processes and reduce the uncertainties associated with the existing measurements.

The increased luminosities over the years were also accompanied by higher pileup, as shown in Fig. 3.4. Pileup complicates event reconstruction and data analysis, as the several collision events must be disentangled from one another. To better manage the high PU, the detectors have undergone significant upgrades for better resolution and implemented sophisticated algorithms for precise event reconstruction. Modifications have also been made to the trigger systems to improve the online data selection. Further discussions on the improvements made by the CMS experiment are detailed in Sec. 3.2 and Sec. 4.2.

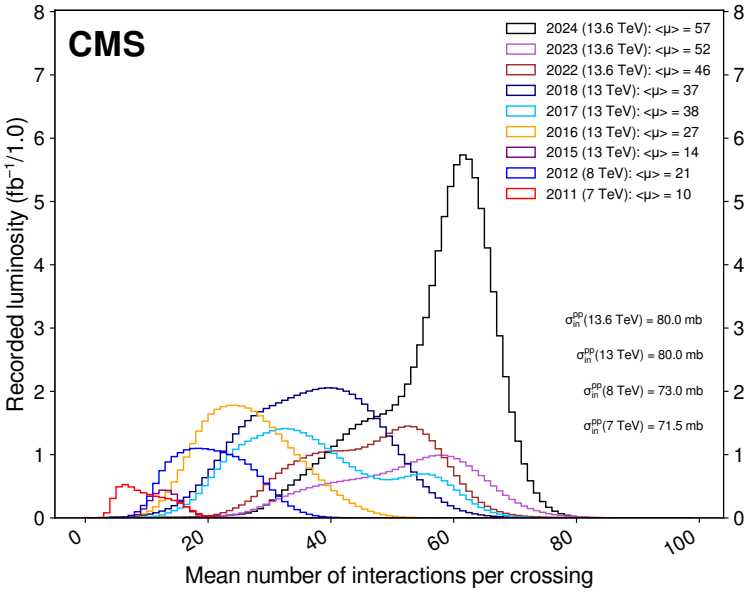

Figure 3.4: Distribution of the mean number of interactions per bunch crossing (pileup) for pp collisions across data taking years. The overall mean values per year and the minimum bias cross sections corresponding to each center-of-mass energy are also shown [84].

The end of Run 3 will also conclude the Phase 1 operations at the LHC. Subsequently, the LHC will enter Long Shutdown 3 (LS3), a crucial upgrade phase that will extend into early 2030, paving the way for the High-Luminosity LHC (HL-LHC) and Phase 2 operations. The HL-LHC aims to operate at the design center-of-mass energy of 14 TeV and significantly increase the luminosity by 5 to 7.5 times the design value [87]. Cutting-edge technologies like new Niobium-Tin (Nb_3Sn) superconducting magnets that can achieve up to 12 T magnetic fields will be incorporated for sharper beam focusing. Additionally, novel superconducting "crab cavities" will be used at the big experiments to enlarge the overlap area of the two colliding proton bunches, thereby raising the collision probabilities. These advancements also necessitate major updates to detectors and computing systems to manage the surge in data and radiation [88–91]. Efforts towards these upgrades have been underway since as early as LS2. Over its projected 12-year operation, the HL-LHC is expected to deliver between 3000 to $4000\,\text{fb}^{-1}$ of data, enabling deeper studies into rare phenomena and potentially uncovering new physics.

With the overview of the LHC now complete, the next section is dedicated to a detailed description of the CMS experiment which is central to the research presented in this thesis.

3.2 The Compact Muon Solenoid Experiment

The Compact Muon Solenoid detector is one of the two large multipurpose detectors at the LHC [76]. Located 100 m underground in an experimental cavern at IP5 near the French village of Cessy, this detector was designed to investigate a broad range of physics at the TeV scale, including the search for the Higgs boson(s), extra dimensions, dark matter, precision measurements of SM particles and interactions, flavor physics, and heavy-ion physics.

The CMS experiment takes its name from the unique characteristics of its detector design. The term compact stems from the relatively small yet dense structure of the detector. Measuring 21 m in length and 15 m in diameter and weighing 14,000 tonnes, the detector is six times smaller in volume compared to ATLAS but twice as heavy. Another distinguishing feature of the experiment is its exceptional capabilities in detecting muons. A significant portion of the detector accommodates the muon chambers, accounting for about 80% of its volume. The CMS detector also integrates one of the most powerful solenoid superconducting magnets ever constructed for a particle physics experiment, achieving a very high magnetic field of 3.8 T, which is essential for accurate momentum measurements within the confined space.

This section provides a detailed description of the current detector system after a quick introduction of the CMS coordinate system in Sec. 3.2.1. A brief overview of the upgrades planned for the detector toward the HL-LHC phase is also given.

3.2.1 Coordinate system

The CMS experiment adopts a right-handed Cartesian coordinate system centered on the nominal collision point. The x-axis points toward the center of the LHC ring, and the y-axis points upward, perpendicular to the plane of the LHC ring. The z-axis aligns with the LHC beamline, pointing along the direction of the anti-clockwise proton beam. Given the cylindrical symmetry of the CMS detector around the beamline, a cylindrical coordinate system is also utilized where r is the radial distance measured from the nominal IP and the azimuthal angle ϕ is measured from the x-axis in the x-y plane. The polar angle θ is measured from the positive z-axis. An illustration of the CMS coordinate system is shown in Fig. 3.5.

While the above-mentioned coordinate systems are well suited for representing the detector, the same cannot be said in the context of analyzing collisions. Their use is complicated by the fact that protons are composite particles made up of partons whose momentum fractions are unknown. This uncertainty makes it difficult to determine the longitudinal boost of the collision's rest frame along the z-axis. To tackle this challenge, physicists rely on Lorentz-invariant variables. These variables, defined in the plane perpendicular to the beam direction, remain unaffected by variations in the longitudinal boost, thereby providing a consistent framework for analyzing collision events. The simplest Lorentz-invariant variables include the transverse momentum p_{T} and the transverse energy E_{T}, defined as

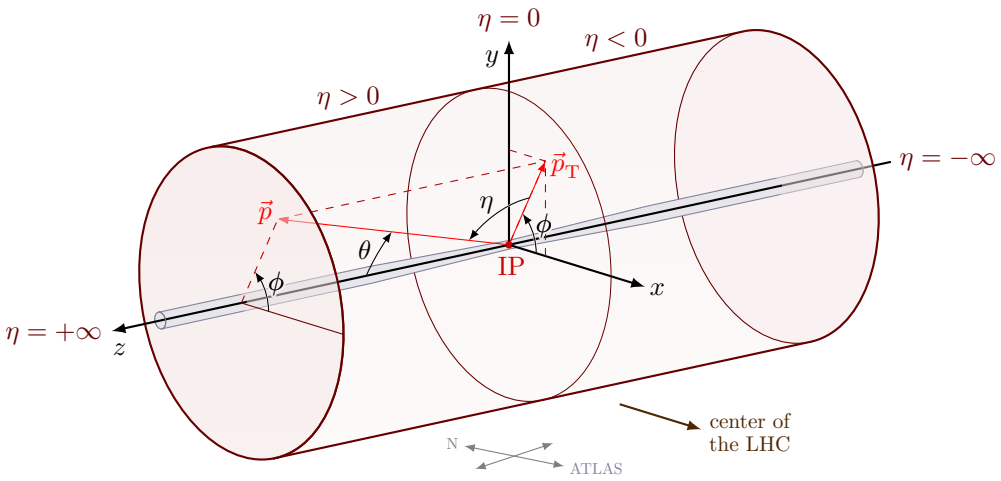

Figure 3.5: Schematic representation of the coordinate system used by the CMS detector. Adapted from [92].

$$p_{\mathrm{T}}^2 = p_x^2 + p_y^2 \tag{3.6}$$

$$E_{\mathrm{T}}^2 = m^2 + p_{\mathrm{T}}^2 \tag{3.7}$$

where p_x and p_y represent the momentum in the x and y directions respectively, and m represents the invariant mass.

The pseudorapidity η, defined as

$$\eta = -\ln \tan \frac{\theta}{2} \tag{3.8}$$

is commonly used instead of θ to describe the angle of a particle relative to the beam axis. Importantly, the difference in pseudorapidity, $\Delta\eta$, between two particles is Lorentz invariant.

The separation in the η-ϕ plane between particles, commonly represented as ΔR, where

$$\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}, \tag{3.9}$$

forms another critical Lorentz invariant metric.

3.2.2 Detector components

The CMS detector is designed as a symmetric cylindrical structure centered around the beam line. The central part of the detector, known as the barrel, surrounds the collision point. It is designed to capture particles that are ejected perpendicular to the beam line. The endcaps cover the ends of the detector and are crucial for detecting particles that travel along the direction of the beam close to the beam line. Together, the barrel and endcaps provide comprehensive spatial coverage.

In order to accurately measure a wide range of particle interactions and detect different signatures and final states, the CMS experiment incorporates several concentric subdetectors. These include the innermost silicon tracker registering charged particle trajectories, the electromagnetic and hadron calorimeters measuring particle energies, and the muon system tracking muons. Positioned between the calorimeters and the muon system, the superconducting solenoid generates the powerful magnetic field required for bending the paths of charged particles and measuring their momentum. Figure 3.6 presents the overall layout of the CMS detector, highlighting its different layers.

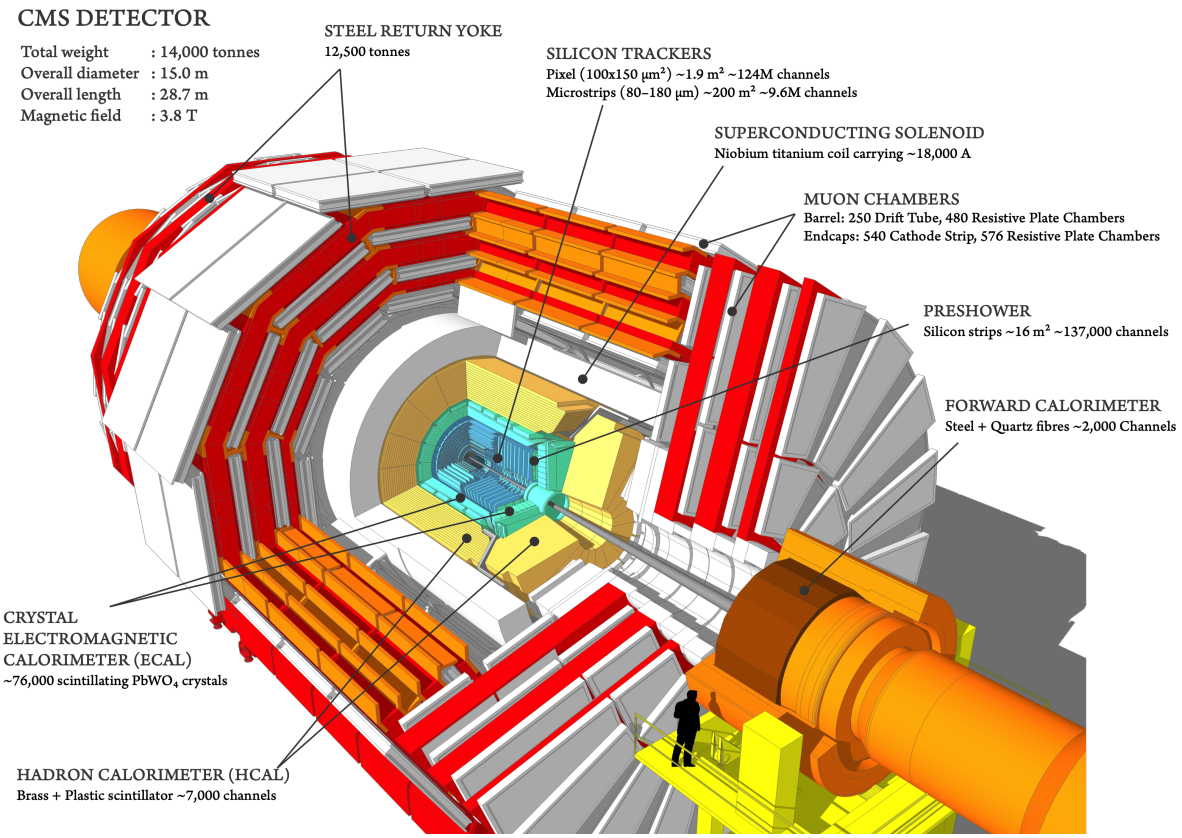

Figure 3.6: Schematic drawing of the CMS detector [93].

The following text provides an in-depth description of each component of the CMS detector.

3.2.2.1 Superconducting solenoid magnet

The magnet is the core component of the CMS detector, designed to provide a strong magnetic field of 3.8 T, which is approximately 100,000 times stronger than the Earth's magnetic field. This intense field bends the trajectory of charged particles, thus allowing precise determination of the particles' energy, momentum, and charge as they pass through the detector. The solenoid, one of the largest ever built, has a total length of 12.5 m, with an inner diameter of 6 m. It operates with an electric current of 16.2 kA. The total energy stored in the magnet at full power is approximately 2.3 GJ. The solenoid's coils are constructed with NbTi superconducting wires, coated with aluminium for stability under high current loads.

To maintain superconductivity, these coils are cooled to 4.5 K (-268.3° C) by a sophisticated cryogenic system that uses liquid helium as the primary coolant.

With the exception of the muon chambers, all the detector systems are located within the solenoid, where they are exposed to the strong, uniform magnetic field necessary for accurate particle measurements. Surrounding the solenoid is a 12,500-tonne iron yoke that serves to return the magnetic flux generated by the superconducting coil while also providing structural support to the whole detector. With a diameter of 14 m and a length of 21.6 m, the yoke also houses the muon detectors which are interleaved within its layers. The yoke generates a magnetic field of approximately 1.8 T that enhances muon detection by ensuring accurate tracking of particles as they pass through the outer regions of the detector.

3.2.2.2 Silicon Tracker

Positioned only 2.9 cm away from the IP, the silicon tracker is the innermost component of CMS [94–96]. It provides high-precision measurements of hits produced in its sensors by the passing charged particles, which are combined to reconstruct particle trajectories or tracks. The curvature of these tracks, induced by the CMS magnetic field, helps determine particle momentum p_{T} as:

$$p_{\mathrm{T}}\,[\mathrm{GeV}] = 0.3 \cdot B\,[\mathrm{T}] \cdot R\,[\mathrm{m}], \tag{3.10}$$

where B and R represent the magnetic field strength and the radius of curvature, respectively. The tracks are also traced back to identify the original hard scatter interaction points or the primary vertices. Additionally, the tracker identifies secondary vertices and displaced tracks, often resulting from the decay of long-lived particles like heavy-flavor hadrons or τ leptons. This displacement data is essential for flavor tagging, especially b-tagging, which is crucial for identifying jets from heavy-flavor quarks.

The tracker occupies a cylindrical volume, 5.8 m in length and 2.5 m in diameter, which is permeated by a uniform 3.8 T magnetic field. It provides pseudorapidity coverage in the region $|\eta| < 2.5$. It is composed entirely of silicon-based sensors chosen for their high spatial resolution, rapid response time, and excellent radiation hardness—all qualities particularly suited to the demanding environment near the IP. With over 200 m^2 of active silicon area, the CMS tracker is the largest silicon tracker ever constructed.

The tracker is divided into two main subdetectors: the inner pixel detector, which provides detailed imaging close to the collision point, and the outer strip detector, which extends the tracking capabilities outward. Figure 3.7 presents a schematic representation of one-quarter of the tracker, illustrating the arrangement of pixel and strip detector layers along the r and z dimensions.

Initially, during Run 1 and the 2016 data-taking, the tracker used the so-called Phase-0 pixel detector, designed for nominal LHC conditions. This detector consisted of three barrel layers (BPIX) and two forward/backward disks (FPIX). However, as the LHC's instantaneous luminosity substantially exceeded initial projections, it became necessary to upgrade the pixel detector to handle the increased collision rates while maintaining decent tracking performance. The 2016 year-end technical stop (YETS) marked the transition to the Phase-1 pixel

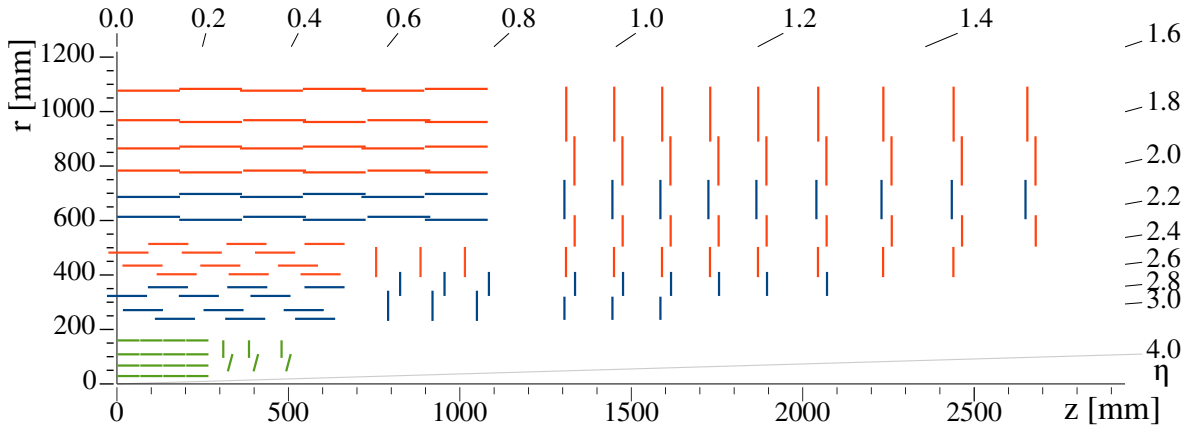

Figure 3.7: An $r-z$ view of one-quarter of the silicon tracker. The layers of the pixel detector are drawn as green lines. The red and blue lines represent the strip detector layers [97].

detector, which features an enhanced configuration of four barrel layers and three endcap disks [95, 96]. Figure 3.8 provides a comparison of the layouts of the Phase-1 and Phase-0 pixel detectors.

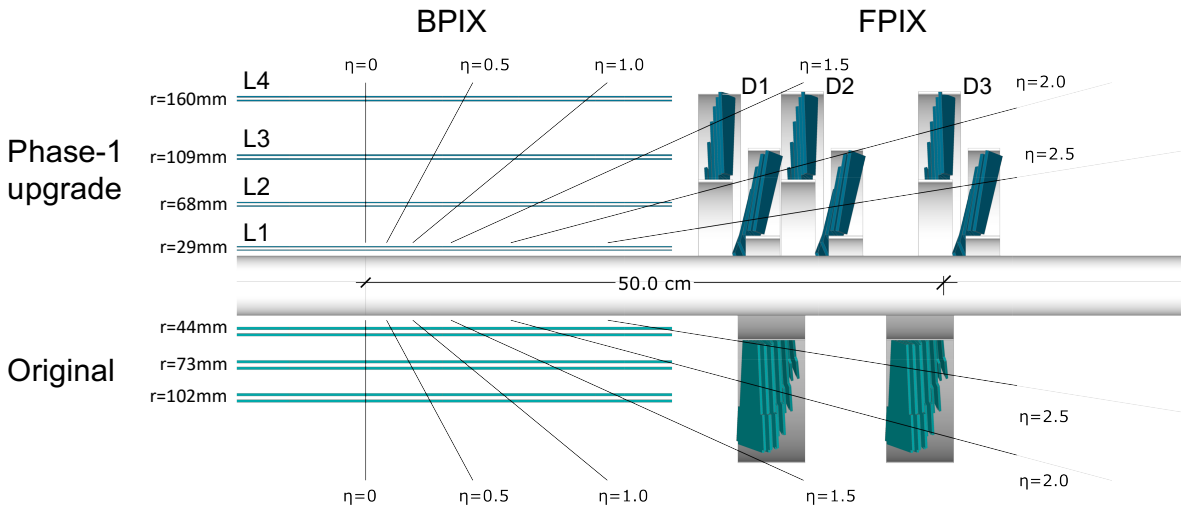

Figure 3.8: Comparison of the longitudinal view of the upgraded Phase-1 pixel detector and the original Phase-0 detector [96].

In the Phase-1 pixel detector, the barrel layers are positioned at radii of 2.9 cm, 6.8 cm, 10.9 cm, and 16.0 cm, while the layers of the FPIX are located at ± 29.1 cm, ±39.6 cm and ±51.6 cm from the IP. The detector occupies a total active silicon area of 1.9 m^2 and consists of 1856 segmented silicon sensor modules—1184 in the BPIX and 672 in the FPIX. Employing n-in-n technology, the sensor in every module includes 160 × 416 pixels, with each pixel measuring approximately 100 × 150 microns. Each sensor is bump-bonded to 16 readout chips (ROCs). Altogether, this complex array comprises 124 million readout channels.

Notably, the upgraded detector features an innermost layer that is located much closer (at $r = 2.9$ cm) to the collision point as compared to that of the Phase-0 detector (at $r = 4.4$ cm). This placement was facilitated by the installation of a new narrower beam pipe with a radius

of 2.3 cm as opposed to the original 3.0 cm. Moreover, the addition of a fourth barrel layer at $r = 16.0$ cm and two endcap disks at $|z| = 51.6$ cm ensures a four-hit coverage over the entire η range. The new components provide redundancy in pattern recognition and reduce fake rates in high pileup conditions. They also reduce the distance between the last pixel hit and the strip detector, thereby improving track extrapolation towards the strip detector.

Significant attention was also devoted during the upgrade to managing the material budget, as minimizing the material in the tracker is crucial to prevent multiple scattering and nuclear interactions that could compromise tracking. Major improvements included the use of lightweight carbon-fiber materials for the detector's support structure and a lower-mass, two-phase CO_2 cooling system. The electronic boards were also strategically relocated to regions of higher pseudorapidity, outside the primary tracking zone. As a result, despite the addition of extra sensor layers to the Phase-1 detector, the material budget in the central region remained largely unchanged from the Phase-0 detector, while significant reductions were achieved in the forward region at $|\eta| > 1$. Figure 3.9 displays the material budget of the Phase-1 pixel detector compared to the Phase-0 pixel detector.

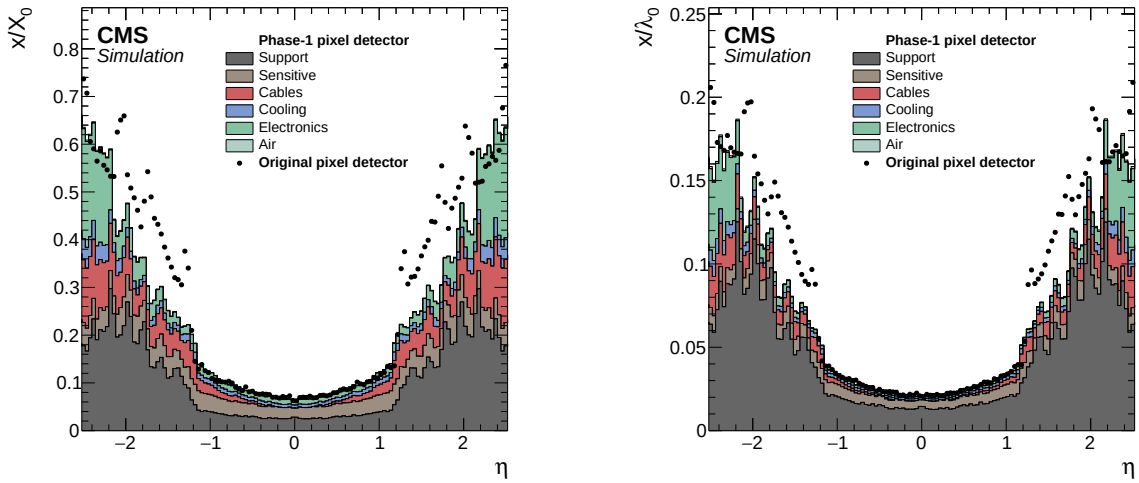

Figure 3.9: Material budget in the pixel detector in units of radiation lengths X_0 (left) and hadronic interaction lengths λ_0 (right) as a function of pseudorapidity, as obtained from simulation. The material budget of the original pixel detector (black dots) is compared to the Phase-1 pixel detector (stacked histograms) within the tracking acceptance of $|\eta| < 2.5$. A considerable reduction of the material budget is obtained following the Phase-1 upgrade [96].

During LS2, all the modules of BPIX layer 1, which are the most irradiation-exposed, were completely replaced to ensure optimal performance in the higher luminosities expected in Run 3 [93]. Since the particle hit rates, and hence the damage to the sensors, in the other BPIX layers and the FPIX disks are considerably lower, no further replacements were deemed necessary in these areas.

The intricate design of the pixel detector significantly contributes to its operational capabilities. The detector achieves a spatial resolution of approximately 10 μm in the r-ϕ direction and about 20 μm in the z direction [98]. This precision is facilitated by the small dimensions of the pixels, which allow for maintaining a low occupancy despite the high particle flux

near the IP. Additionally, the resolution is enhanced by the charge sharing induced by the magnetic field among neighboring pixels.

As the particle flux decreases rapidly with the radial distance (following an inverse square relationship), the detector occupancy decreases accordingly. This allows for lower granularity in the outer parts of the tracker, further from the IP. So, the outer strip detector is composed of micro-strip sensors. This detector encompasses a cylindrical volume of 6 m in length and 2.2 m in diameter around the beam pipe. It consists of 15 148 individual silicon modules containing 9.3 million readout strips. Each module features a p-in-n type silicon sensor with a uniform n++ backside implant.

The barrel section of the strip detector comprises the Tracker Inner Barrel (TIB) and the Tracker Outer Barrel (TOB), containing four and six concentric layers of silicon modules, respectively. Located on either side of the barrel, are the two small Tracker Inner Disks (TID) and the larger Tracker Endcaps (TEC). The TID consists of three disks, each with three rings of modules. Each of the two TECs comprises nine disks, with the number of rings decreasing from seven in the initial three disks to four in the ninth disk. This gradation ensures coverage for $|\eta| < 2.5$ while minimizing silicon usage in low occupancy areas.

The dimensions of the strip sensors vary by location to optimize signal-to-noise ratios: sensors in the inner parts of the strip detector are 320 μm thick, while those in less irradiated, lower occupancy areas are 500 μm thick. The pitch of the sensors, crucial for resolving closely spaced hits, ranges from 80 μm to 205 μm, increasing with radial distance. This variation achieves finer spatial resolution where particle density is high and coarser resolution where tracks are more dispersed, thereby balancing precision with efficiency.

The modules in the barrel are oriented parallel to the beam line, facilitating precise measurement of the r-ϕ coordinates. The wedge-shaped modules in TID and TEC, with radially outward-facing strips, enable measurement in ϕ and z. To allow for the measurement of an additional coordinate (z in the barrel and r in the endcaps), the first two layers of both TIB and TOB, rings 1 and 2 of the TID, and rings 1, 2, and 5 of the TEC employ "stereo" modules, depicted in blue in Fig. 3.7. These consist of two independent silicon modules mounted back-to-back with the sensor strips in the second module tilted by 100 mrad, for three-dimensional point reconstruction.

The strip detector has been in operation since its installation in December 2007. Even after more than ten years of data-taking, it showcases excellent performance. At the time of writing this thesis, approximately 96% of the total modules of the strip detector are active and collecting data. The hit resolution of the detector ranges between 20 and 40 μm, depending on the strip pitch [99]. This performance significantly surpasses the theoretical expectations for binary readout, calculated as pitch/$\sqrt{12}$, underlining the detector's reliability.

The optimal performance of the CMS tracker is crucial for physics analyses, including the one presented in this thesis. Accurate tracking demands precise knowledge of each hit's position and resolution. This precision is achieved through rigorous tracker alignment, which systematically corrects misalignments and mechanical shifts within the tracker system that could arise from environmental or operational stresses [100, 101]. By adjusting the geometric parameters of the tracking elements to align with a global reference frame, the alignment ensures that spatial coordinates accurately capture the true trajectories of particles. Such

precision is vital for detailed vertex reconstruction, momentum analysis, and b-tagging.

During the course of the doctoral research, the author has made significant contributions to the CMS tracker alignment efforts, leading to marked improvements in the tracker's performance. These contributions are highlighted in several publications [12–17]. Additional details on some of these efforts are further discussed in Appendix A of this thesis.

3.2.2.3 Electromagnetic Calorimeter

Placed outside the silicon tracker, the electromagnetic calorimeter (ECAL) is a hermetic, homogenous calorimeter designed to measure the energies of electrons and photons with high accuracy [102]. The ECAL is constructed from lead-tungstate ($PbWO_4$) crystals that scintillate when an electron or photon passes through them. The high density ($8.28\,g/cm^3$), short radiation length (0.89 cm), and small Molière radius (2.2 cm) of lead tungstate enable rapid absorption of the energy of the incoming particles, thereby allowing the construction of a compact yet highly granular calorimeter. Furthermore, the radiation-hard $PbWO_4$ crystals exhibit fast scintillation decay, emitting about 80% of the light within 25 ns, thereby quickly disentangling particles originating from neighboring bunch crossings in high-rate environments of the LHC.

A schematic representation of the ECAL is provided in Fig. 3.10, with a volumetric view in the top panel and the longitudinal view of one quadrant in the bottom panel. The ECAL consists of two main sections: the barrel (EB), covering the pseudorapidity range $|\eta| < 1.479$, and the endcaps (EE) that extend the coverage to $1.479 < |\eta| < 3.0$. The EB consists of 61 200 crystals, shaped as truncated pyramids with a front-face crystal cross section of approximately $2.2 \times 2.2\,cm^2$ and a length of 230 mm (or 25.8 times the radiation length, X_0). The crystals are grouped into 36 *supermodules*. Each supermodule subtends 20° in ϕ and comprises four modules with 500 crystals in the first module and 400 crystals in each of the remaining three modules. Each ECAL endcap disk holds 7324 tapered right-sided crystals—each 22 cm long (or 24.7 X_0) with a front-face cross section of $2.86 \times 2.86\,cm^2$. The crystals in each endcap are housed in two semicircular structures called *dees* and are grouped into mechanical units of 5×5 crystals called *supercrystals*. In both the EB and EE, the crystals are arranged in a quasi-projective geometry, with their axes tilted up to 3°radially to minimize energy losses at crystal boundaries from shower sharing and to avoid gaps that might allow particle energy to escape. The scintillation light produced by the EB and EE crystals is detected using avalanche photodiodes (APDs) and vacuum phototriodes (VPTs), respectively, selected for their speed, radiation tolerance, and ability to operate in the 3.8 T magnetic field.

As depicted in Fig. 3.10, the EB and EE are complemented by the Preshower (ES) detector, positioned in the fiducial region $1.653 < |\eta| < 2.6$. Constructed as a sampling calorimeter, the ES features two main components: lead radiator plates that initiate electromagnetic showers, interspersed with planes of 2 mm-thick silicon strip sensors that measure the deposited energy and map the transverse shower profiles. The ES is designed to distinguish high-energy forward photons from double-photon signals of $\pi^0 \to \gamma\gamma$ decays. It also provides fine-grained spatial resolution, which is crucial for accurately determining the position and the incident angle of the particles entering the ECAL.

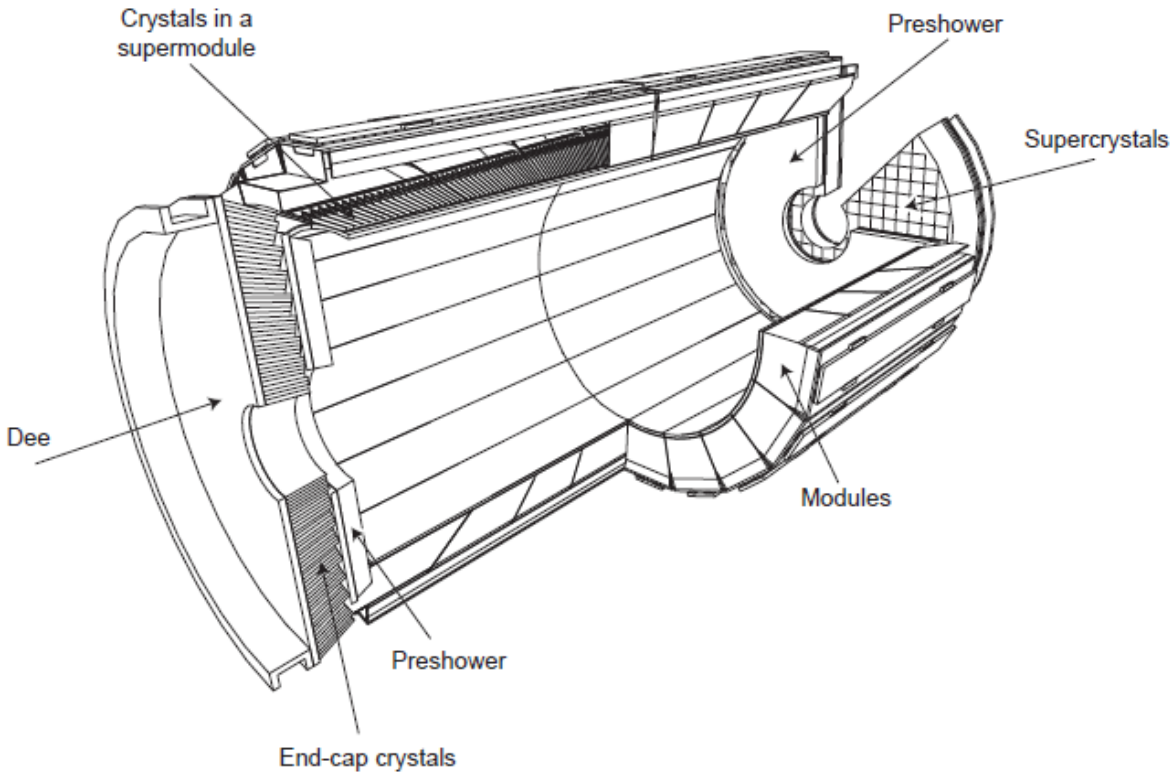

(a) A three-dimensional view of the electromagnetic calorimeter [76].

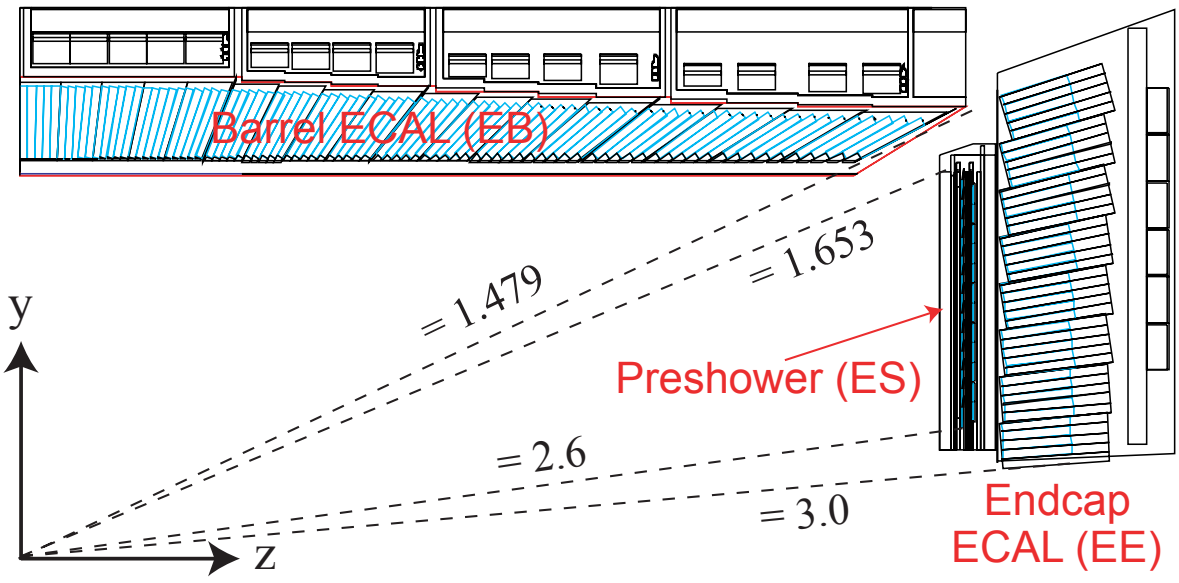

(b) Schematic longitudinal view of the layout of a quarter of the electromagnetic calorimeter. The ECAL Barrel (EB), the Endcap Preshower (ES), and the ECAL Endcap (EE) are shown, and the pseudorapidity envelopes are specified [102].

Figure 3.10: Layout of the CMS electromagnetic calorimeter.

The energy resolution of any homogenous calorimeter can be expressed as:

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c, \tag{3.11}$$

where a, b, and c represent the stochastic, noise, and constant terms, respectively. In the ECAL, a arises from statistical fluctuations in the energy deposited within the detector. The noise term b accounts for the electronics, digitization, and pileup noise. The constant term c, which dominates the resolution for high-energy electron and photon showers, includes systematic effects such as calibration errors, energy leakage, non-uniformities in the detector material, and geometrical imperfections. A measurement of the ECAL barrel energy resolution for electrons performed in a test-beam environment yielded values of a, b, and c corresponding to 2.8%, 12%, and 0.3%, respectively [103].

During Run 2 of the LHC, the energy resolution for electrons emanating from Z boson decays and entering the ECAL without significant bremsstrahlung losses was better than 1.8%, 3.0%, and 4.5% in the $|\eta|$ intervals [0.0, 0.8], [0.8, 1.5], [1.5, 2.5], respectively [104].

3.2.2.4 Hadron Calorimeter

The hadron calorimeter (HCAL), installed between the ECAL and the steel yoke, is crucial for measuring the energies of charged and neutral hadrons. Together with the ECAL, it is also used to estimate the missing transverse energy that results from neutrinos and non-interacting, uncharged particles [105]. As hadrons penetrate the HCAL, they initiate a complex series of interactions with the nuclei of the detector material, producing cascade particle showers that include both hadronic and electromagnetic components. To manage these intricate cascades effectively, the HCAL is designed as a highly-segmented sampling calorimeter. It alternates layers of high-density material that stop the cascades with scintillators that convert the deposited energy into measurable light signals.

As illustrated in Fig. 3.11, the HCAL consists of four major subdetectors, each covering distinct fiducial regions: the HCAL Barrel (HB), the HCAL Endcap (HE), the HCAL Outer Barrel (HO), and the HCAL Forward (HF) detectors.

The HB and HE cover the pseudorapidity range $|\eta| < 1.3$ and $1.3 < |\eta| < 3.0$, respectively, and are constructed using brass as the absorber material and plastic scintillators as the active medium. The HB is segmented into 36 identical wedges spanning 20°in azimuth, similar to the EB. Each wedge comprises 14 inner brass plates, alternating with 17 layers of scintillator tiles. The HB has roughly 40 000 scintillator tiles in total. For efficient data readout, the HB features a 72-fold segmentation along ϕ and has 32 divisions in η. This results in 2304 projective units in the η-ϕ space called *towers*, each with a lateral dimension of 0.087×0.087. The HE, on the other hand, consists of 18 wedges, each with 18 scintillator layers. The projective towers in the HE have a segmentation in the η-ϕ space of 0.087×0.087 (0.17×0.17) for $|\eta| < 1.6$ ($|\eta| > 1.6$).

Positioned outside the CMS magnet, the HO enhances the HCAL performance at $|\eta| < 1.3$ by measuring the energy from highly energetic hadronic showers that pass the HB and punch through the solenoid. It uses plastic scintillators and exploits the magnet's iron flux-return

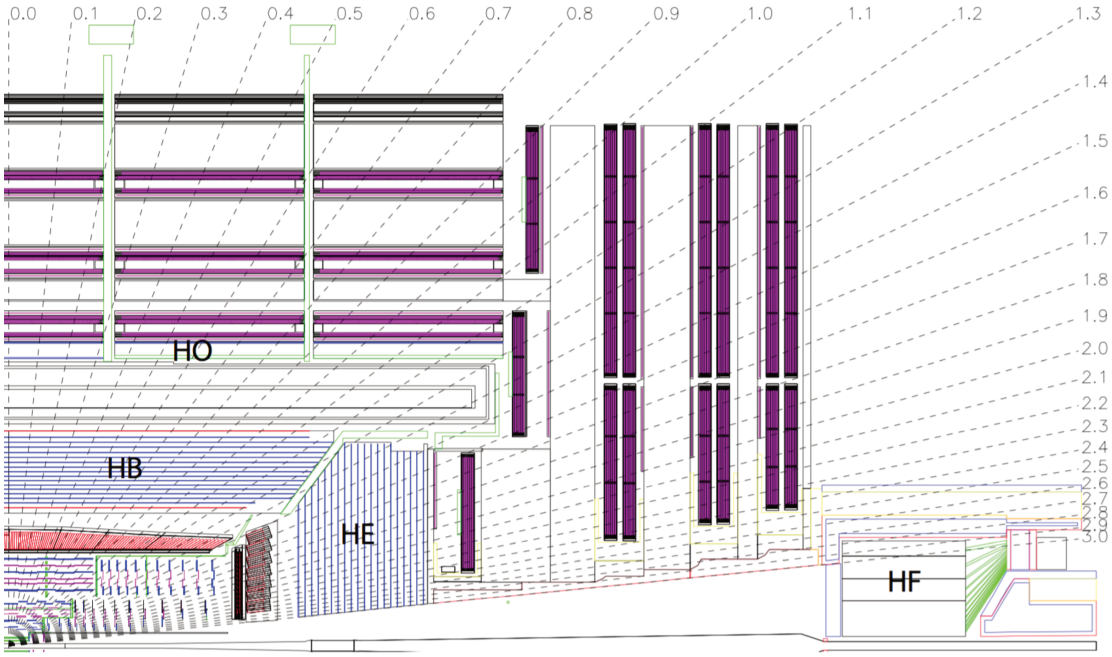

Figure 3.11: Schematic view of a quarter of the HCAL, showing the positions of its four major components: HB, HE, HO, and HF. [105]

yokes as the absorber material. The assembly is divided into 5 rings, each comprising 12 identical ϕ sectors, with an η-ϕ segmentation closely matching that of the HB.

The light generated in the plastic scintillators of HB, HE, and HO is collected by wavelength-shifting (WLS) fibers, summed optically, and channeled to photodetectors for conversion into digital signals. Initially, hybrid photodiodes (HPD) were used due to their tolerance to magnetic fields, high gain, and linear response. However, they were prone to creating high-amplitude noise from electrical discharges and ion feedback. So, as part of the HCAL Phase-1 upgrade, the HPDs in all three subdetectors were replaced by silicon photomultipliers (SiPMs)—in the HO during LS1, in the HE during the 2017-2018 YETS, and in the HB during LS2 [93,106]. SiPMs offer three times the photon detection efficiency and 200 times the gain of HPDs. The improved signal-to-noise ratio of the SiPMs allows for a finer depth[1] segmentation in the HB and HE. Figure 3.12 shows the depth segmentation of the HCAL following the Phase-1 upgrade. Compared to the original design, the HB readout segmentation increased from one to four channels, whereas for HE, it increased from two to three channels per tower to six to seven channels. This upgrade improves the tracking of hadronic showers and mitigates radiation damage, especially in high-η regions.

Lastly, the forward hadron calorimeter, HF, provides coverage in the region $3.0 < |\eta| < 5.2$ and is essential for detecting particles in the very forward regions of the CMS detector. It is a quartz-fiber Cherenkov calorimeter with steel absorbers strategically placed outside the solenoid at approximately $|z| = \pm 11$ m from the IP. Quartz is employed as the active medium

[1] Here, "depth" refers to the layers of scintillator material within a tower that are grouped together and read out simultaneously by a single photodetector.

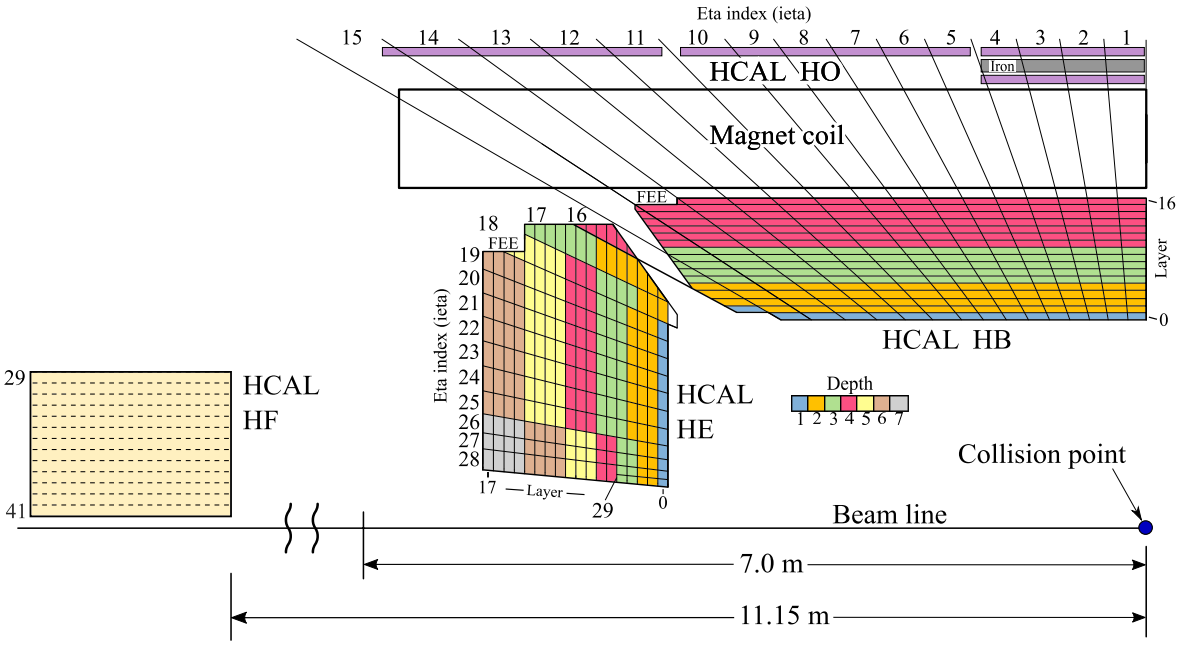

Figure 3.12: Depth segmentation of the HCAL achievable after the Phase-1 upgrade. Within a tower, layers with the same color are routed to the same SiPM [93].

due to its radiation-hard properties, crucial in the forward regions where particle fluxes can be as high as $10^{11}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The detection method in the HF relies on the emission of Cherenkov light by secondary charged particles as they pass through the quartz fibers. The light collected from HF calorimeter fibers is converted to charge by photomultiplier tubes (PMTs). As part of the Phase-1 upgrade, all the PMTs were replaced with newer, more radiation-hard PMT models during LS1 [93, 106].

The relative charged pion energy resolution for the combined ECAL and HCAL systems was estimated in test beam measurements as [107]:

$$\frac{\sigma}{E} = \frac{84.7\%}{\sqrt{E}} \oplus 7.6\%, \tag{3.12}$$

where E is the energy in GeV. Using 13 TeV collision data from 2016, the pion energy resolution was found to be 19.4% in HB, 18.8% in HE, and 23.6% in the transition regions [108].

3.2.2.5 Muon system

As evident from the very name of the experiment, accurate muon detection is of central importance to CMS, with the muons playing a pivotal role in many physics results, including the discovery of the Higgs boson. The CMS muon system, designed to identify muons, measure their momenta, and effectively trigger on them, is strategically placed outside the other detector components, within the flux-return yoke of the solenoid [76, 109]. Since high-energy muons travel several meters before decaying and interact only weakly with matter, they can typically pass through the dense materials of the calorimeters, solenoid, and iron yoke with minimal energy loss.

The muon system is designed for precise and robust measurement of muons across the entire kinematic range of the LHC, necessitating a setup with high redundancy and expansive surface area coverage. Because the muon system consists of about 25 000 m^2 of detection planes, the muon chambers needed to be inexpensive as well as robust. So, gaseous detectors, that operate by capturing the ionization produced when charged particles pass through a gas-filled chamber, were the choice. Their reliability in high-radiation environments makes them an ideal choice for ensuring long-term performance under the extreme conditions of the LHC.

The muon system currently comprises three main subdetectors: the drift tubes (DTs), the cathode strip chambers (CSCs), and the resistive plate chambers (RPCs). In anticipation of the HL-LHC phase, the system is being enhanced with the addition of gas electron multipliers (GEMs) [110]. The first of these GEM chambers was integrated during LS2 and is currently operational for Run 3. The physical arrangement of these muon detectors is shown in Fig. 3.13. The central section of the muon system features a barrel configuration with DTs and barrel RPCs whereas the endcap section comprises CSCs, endcap RPCs, and the newly installed GEMs.

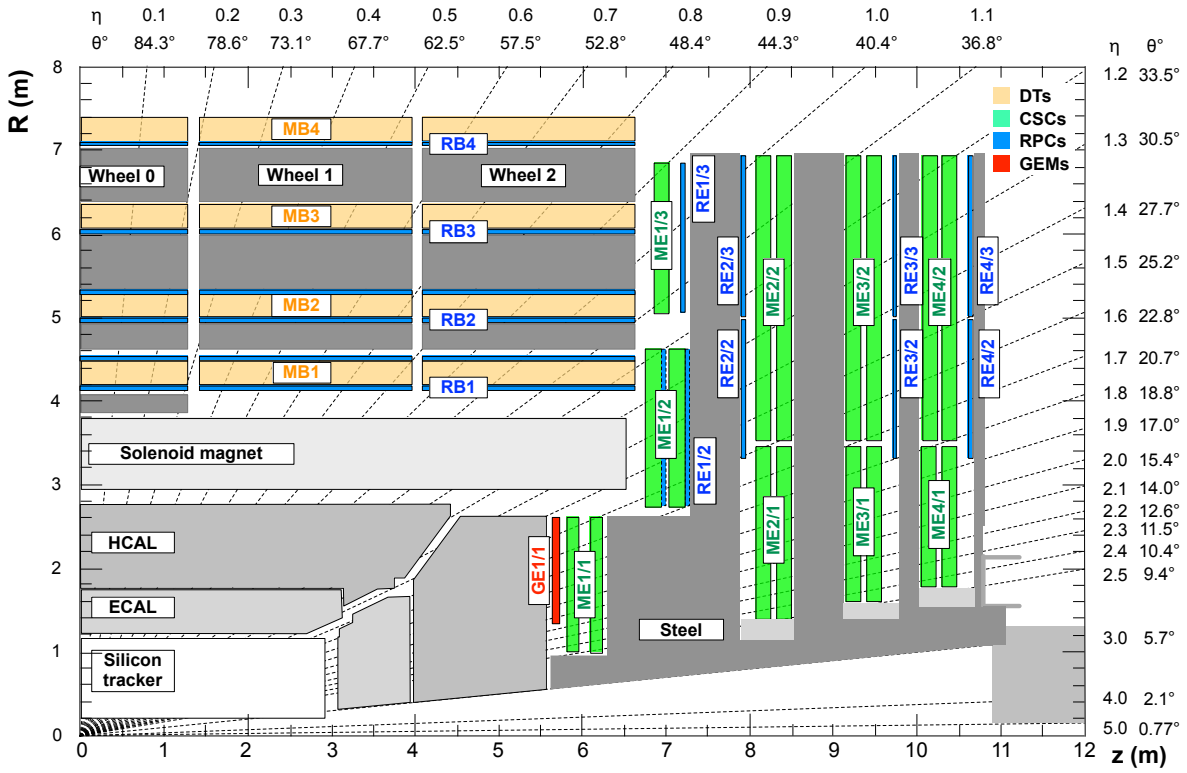

Figure 3.13: Schematic view of a quadrant of the CMS muon system in the r-z plane. The different components, drift tubes (DTs, yellow), cathode strip chambers (CSCs, green), resistive plate chambers (RPCs, blue), and gas electron multipliers (GEMs, red), are shown. The iron yoke is illustrated by the dark gray areas [111].

In the region with $|\eta| < 1.2$, characterized by a low muon rate and a uniformly distributed magnetic field, muon detection is handled by the drift tubes. The basic DT units are rectangular drift cells measuring $4.2 \times 1.3\,cm^2$. The cells contain a gas mixture of 85% Ar, providing efficient ionization, and 15% CO_2 acting as a quenching agent. A gold-plated stainless steel

central anode wire is stretched along the length of each cell, with aluminum cathode sheets on its side walls, which generate the electric field needed for the drift process. Electrode strips located at the top and the bottom help shape the drift field. The DT system is composed of 250 chambers, organized into four stations—labeled MB1-MB4—wedged between the layers of the flux-return yoke at consistent radial distances. Each chamber consists of eight to twelve staggered layers of drift cells. Some of these layers are oriented to measure the muon position in the r-ϕ plane, whereas others provide measurements in the r-z plane. The DT barrel stations are also segmented into five *wheels* along the z dimension, with 12 ϕ-sectors per wheel. With spatial and time resolutions of approximately 100 μm, and 2 ns, respectively, the DT chambers facilitate precise tracking and triggering capabilities.

Cathode strip chambers are employed within $0.9 < |\eta| < 2.4$, where the muon flux is high and the magnetic field is large and non-uniform. These multiwire proportional chambers are filled with a gas mixture consisting of 40% Ar, 50% CO_2, and 10% CF_4. Each CSC module comprises six gas layers, with each layer featuring a plane of radial cathode strips and a plane of anode wires running perpendicular to the strips. This configuration allows the determination of two positional coordinates—azimuthal and radial—for each passing particle. A total of 540 trapezoidal CSC modules are organized into four stations—labeled ME1-ME4—between the disks of of the iron yoke. The stations are further segmented in the radial direction into two or three *rings*, wherein each CSC chamber subtends a 10°or 20°angle in ϕ. A high level of segmentation enables the CSCs to achieve time and spatial resolution of 3 ns and 50-140 μm, respectively, even under high particle rates.

In the region $|\eta| < 1.9$, the DT and CSC are complemented by resistive plate chambers. These gaseous parallel-plate detectors contain two 2 mm wide gas gaps with a copper readout plane in between. They are filled with a gas mixture of 95.2% $C_2H_2F_4$, 4.5% i-C_4H_{10}, and 0.3% SF_6. High-pressure laminate (HPL) plates, coated with conductive graphite, form electrodes, which are read out by aluminium strips. In the barrel region, six layers of RPCs are integrated into each of the five wheels that house the DT chambers. The initial four layers are positioned on the inner and outer sides of the first two DT stations, with the final two layers located on the inner side of the third and fourth stations of the DTs. In the endcaps, four RPC stations are interleaved with the CSC chambers in these stations. With a total of 1056 chambers (480 in the barrel and 576 in the endcaps), the RPC system delivers a rapid response with a time resolution of 1.5 ns, effectively handling high event rates and unambiguously identifying the bunch crossing corresponding to a muon trigger candidate.

During LHC Run 2, the muon system achieved excellent hit reconstruction efficiency of 94-99%, muon reconstruction and identification efficiency of over 96%, and differentiated between prompt muons and those from weak decays within jets with an efficiency higher than 95% [111].

During LS2, new gas electron multiplier detectors were installed in the region covering $1.55 < |\eta| < 2.18$ [93]. This station denoted GE1/1, is the first of three GEM rings planned for the HL-LHC, aimed at managing the very high hit rate expected in the forward region. The GEM detectors are made of three layers of 50 μm thick copper-cladded polyamide foil, with four gas gaps containing an Ar+CO_2 (70:30) gas mixture. Positioned just in front of the first CSC station, each endcap GE1/1 detector features 36 *superchambers*, with each superchamber containing two stacked GEM detectors covering 10°sectors in azimuth. With spatial and

timing resolutions of approximately 100 μm and 10 ns, respectively, the GEM technology can withstand particle rates as high as 1 MHz/cm^2. The GE1/1 chambers have been successfully operated so far during Run 3. An average chamber efficiency of 94% was achieved during 2024 [112]. At the same time, tuning and optimization studies are being done to further enhance the performance.

3.2.2.6 Trigger System

The LHC collides proton bunches in CMS at a maximum rate of 40 MHz. At integrated luminosities exceeding $2 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, on average more than 50 simultaneous proton-proton interactions occur per bunch crossing, leading to a tremendous throughput of data. Given the impracticality and incapability of reading and storing the complete dataset from each subdetector for every bunch crossing, a selective approach is necessary. A significant fraction of these collision events results in low-energy interactions that are not pertinent to CMS's physics objectives, which focus rather on high-energy hard scattering processes. To efficiently filter out these less relevant events and utilize resources to capture those with potential scientific value, CMS utilizes a trigger system that strategically reduces the data acquisition rate by at least a factor of 10^5 [113]. The CMS trigger system is two-tiered, with the first stage of selection performed by the Level-1 (L1) trigger and the second stage by the High-Level trigger (HLT).

The L1 trigger uses custom hardware processors to select up to 110 kHz of the most interesting events[2] utilizing inputs, or Trigger Primitives (TPs), from the calorimeters and muon systems. As depicted in Fig. 3.14, the system comprises multiple layers and components that work in unison to efficiently process and filter collision data [114]. The calorimeter trigger operates in two layers: the first layer processes and sorts energy deposits from ECAL and the HCAL, while the second reconstructs physics objects like electrons, tau leptons, jets, and energy sums, using a time-multiplexed design for efficiency. Simultaneously, the muon trigger system processes signals from CSC, DT, and RPC detectors utilizing three muon track finders, which collectively feed data into the Global Muon Trigger. Both systems then converge at the Global Trigger. Here, a decision is made based on kinematic criteria and a predefined set of algorithms, which, if positive, leads to the readout of full detector data by the data acquisition system (DAQ) for further filtering in the HLT. The entire decision-making process is completed within a latency of 4 μs.

The HLT is a software-based system that refines the selection of events made by the L1 trigger, using data from all subdetectors. It further reduces the nominal event acceptance rate to approximately 1 to 2 kHz. It employs a series of algorithmic steps known as HLT paths, which are executed in a predefined sequence to reconstruct physics objects and implement physics-based selections. Each HLT path starts only after specific Level-1 trigger conditions have been satisfied. Together, these paths constitute a *trigger menu* that determines which data are stored, based on specific physics criteria designed to align with the experiment's objectives. In 2022, the HLT was configured with more than 600 paths to manage proton-proton collision data, with the menu regularly updated to adapt to the dynamic collision conditions at the LHC [93].

[2]The L1 output rate was limited to 100 kHz during LHC Run 2 but was increased to 110 kHz for Run 3.

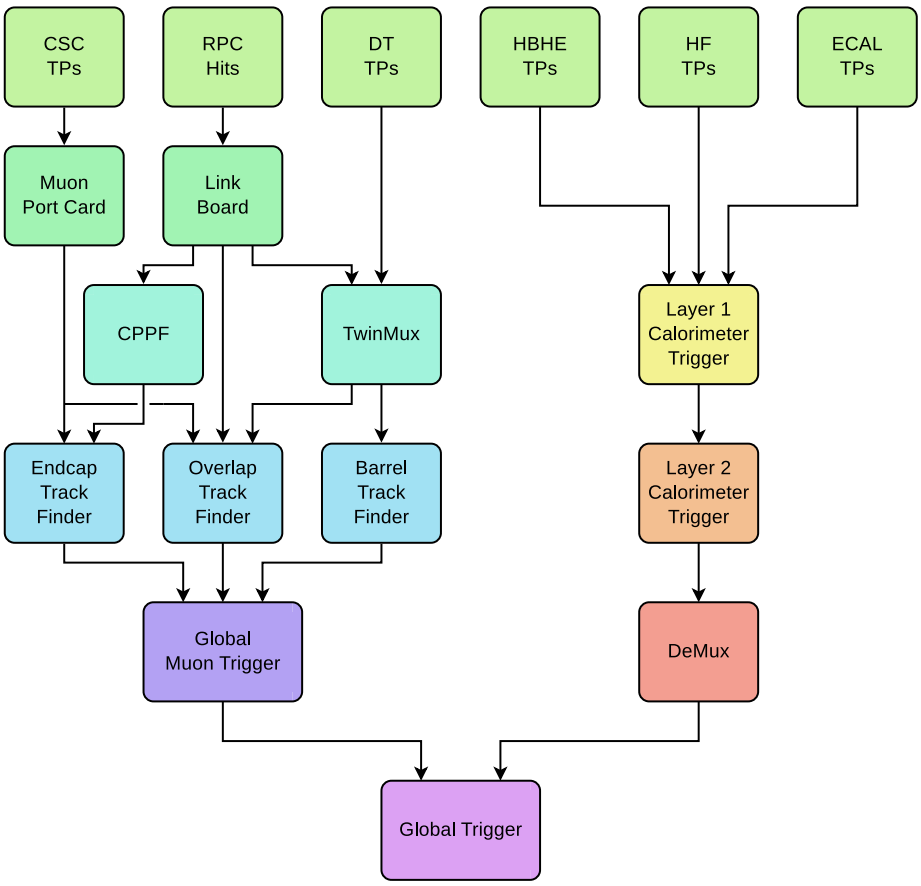

Figure 3.14: Diagram showing the architecture of the L1 trigger system in place during LHC Run 2 and Run 3 [114].

To enhance data analysis capacity beyond standard HLT constraints, CMS has been using parking and scouting techniques since Run 1 [115]. Parking temporarily stores additional raw data, delaying offline reconstruction until computing resources are more available, e.g. during technical shutdowns. This strategy allowed up to 5 kHz additional data in 2018. Scouting, on the other hand, compresses data per event by storing only crucial particle properties like energy, velocity, etc. In 2022, scouting data was recorded at a rate up to 30 kHz at an average size of 13 kB per event, significantly less than the 1 MB size of full raw data.

The events selected by the HLT are then grouped into different *primary datasets* based on the HLT paths that accepted them and stored in the CMS computing facility at CERN, called Tier 0, for further use.

Both L1 and HLT have undergone significant upgrades during Phase 1. The L1 trigger hardware underwent a complete overhaul during LS1. For Run 3, it now incorporates new machine learning techniques and enhanced triggers for detecting long-lived particles, significantly broadening its scientific capabilities. The HLT software was optimized for Run 3 to leverage heterogeneous computing architectures, utilizing both GPUs and CPUs to significantly improve data processing efficiency.

3.2.3 Phase 2 upgrade of the CMS detector

As explained in Sec 3.1.4, the Phase 2 operations of the LHC will deliver unprecedented luminosities of 5 to 7.5 $\times 10^{34}$ $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ to the CMS experiment. With the vast amounts of data recorded, the CMS experiment aims to further the understanding of the Higgs boson by precisely measuring its couplings and examining rare decays [89]. Processes like $\mathrm{H} \to \mathrm{c\bar{c}}$, di-Higgs production, and longitudinal vector boson scattering will become accessible. The searches for new physics phenomena, including supersymmetry, dark matter, and heavy gauge bosons, will also be extended to increase the sensitivity to detect exotic particles. The high luminosity environment of the HL-LHC, however, comes with significant experimental challenges, primarily radiation damage and a high pileup (140 to 200) associated with such high event rates. To address these challenges, extensive upgrades to the CMS detector are necessary to preserve its efficiency, resolution, and background rejection capabilities, ensuring sustained physics performance under these harsh conditions. The upgrades planned for the different components of the detector are outlined in the following.

The upgrade of the silicon tracker involves the complete replacement of the Phase 1 tracker [97]. The new tracker consists of an Inner Tracker (IT) comprised of silicon pixel modules and an Outer Tracker (OT) made of silicon strip and macro-pixel sensors. The IT includes four layers/disks, employing sensors smaller than those used in Phase 1, while the OT has been expanded to six layers. The Phase 2 tracker extends the pseudorapidity coverage to $|\eta| < 4$ and is engineered to withstand fluences up to $2.3 \times 10^{16}\,\mathrm{neq/\,cm^2}$. Increased granularity of the tracker ensures that occupancy remains below 1% across all regions, which is critical for coping with higher pileup levels. A novel module design in the OT facilitates the direct use of tracking information at the L1 trigger.

The CMS experiment will utilize a new dedicated timing detector known as the MIP Timing Detector (MTD) to improve the separation of spatially overlapping but temporally distinct events [116]. The MTD is designed to perform 4D tracking by associating precise timing information with reconstructed tracks. This hermetic detector surrounds the entire Outer Tracker and is comprised of two main components: the Barrel Timing Layer (BTL), which uses Lutetium-yttrium oxyorthosilicate (LYSO) crystals and SiPMs to cover the region $|\eta| < 1.45$, and the Endcap Timing Layer (ETL), which employs Low-Gain Avalanche Diodes (LGADs) to extend coverage to $1.6 < |\eta| < 3.0$. Initially, the MTD will provide a timing resolution of approximately 30-40 ps. Over time, this resolution is expected to degrade to about 50-60 ps in the barrel, with minimal degradation in the endcaps.

For HL-LHC operations, the ECAL barrel will retain its $\mathrm{PbWO_4}$ crystals and photodetectors [117]. Upgrades include new front-end and off-detector electronics to meet the revised L1 trigger requirements. The upgraded electronics will allow the exploitation of the information from single crystals as trigger primitives instead of the 5×5 crystals currently. This shift allows for more precise matching of electromagnetic showers to tracks and improved isolation. The Very Front End (VFE) card will also be replaced to offer better timing resolution and noise filtering to suppress anomalous signals in the photodetectors. For the HCAL, the scintillators and fibers in the HB will continue to be used along with the existing front-end electronics equipped with SiPMs [117]. However, the current HB back-end electronics will be upgraded to sustain the increased trigger requirements of the HL-LHC.

The existing endcap calorimeters, the EE and HE, will be replaced by the High Granularity Calorimeter (HGCAL), a new combined electromagnetic and hadronic sampling calorimeter [118]. The HGCAL primarily uses silicon sensors as the active material throughout most of its volume, supplemented by plastic scintillator tiles read out by SiPMs at large distances from the beam line in the hadronic section. It features high segmentation, with 28 layers in the electromagnetic section and 24 in the hadronic section. The electromagnetic part utilizes compact silicon cells, approximately 0.5 to $1\,\mathrm{cm}^2$ in size, while the hadronic section includes a combination of similar silicon cells and larger plastic scintillators, spanning around 4 to $30\,\mathrm{cm}^2$ in size. This sophisticated design not only enables detailed analysis of electromagnetic showers, taus, and jets but also improves the separation and accurate reconstruction of events, even under high pileup conditions.

While the muon chambers are expected to cope with the increased particle rates, the frontend electronics for DTs, CSCs, and RPCs will be replaced with improved versions to increase radiation tolerance, readout speed, and performance. Additionally, the RPC coverage will be extended from $|\eta| = 1.9$ to 2.4, with two new RPC layers added in the innermost rings of stations 3 and 4. As pointed out in Sec. 3.2.2.5, the muon system will be enhanced in the forward region with GEM detectors, with the first station, GE1/1, already installed during LS2. Two additional GEM stations, GE2/1 and ME0, will be added in the endcap region [119]. The GE2/1 detectors will complement the ME2/1 station, covering the pseudorapidity region $1.6 < |\eta| < 2.4$ and will feature a design similar to GE1/1. A few chambers of GE2/1 have been installed in early 2024, with the remainder scheduled for LS3 and the subsequent technical stops. The ME0 system will extend the muon system coverage even further to $2.0 < |\eta| < 2.8$ and will consist of six-layer GEM detector stacks, each covering a 20° angle in ϕ. Installation of the ME0 station is planned for LS3.

The CMS Level-1 trigger upgrade will increase the maximum event selection rate to 750 kHz [120]. The total latency will be increased to $12.5\,\mu\mathrm{s}$ to include the data from the tracker and high-granularity calorimeter at the L1 level for the first time. With the availability of L1 tracks, the use of advanced offline reconstruction techniques, such as particle flow, becomes possible at the L1 trigger. Additionally, this upgrade also introduces a 40 MHz scouting system that will utilize trigger primitives and objects from sub-detectors to optimize physics selectivity and manage data volumes effectively. For the HLT, the objective remains to reduce event rates by a factor of 100 [121]. Heterogeneous computing architecture appears key to improving the HLT system's efficiency and scalability. Starting in 2022, CMS has already begun integrating heterogeneous hardware, such as GPUs, within the HLT farm to develop the necessary expertise for Phase 2 deployment. Ongoing studies continue to refine and assess the feasibility and impact of these advanced computing solutions for handling the expected computational demands at HL-LHC.

In summary, huge R&D efforts are in progress within the CMS Collaboration for the HL-LHC era. All the planned upgrades are currently undergoing extensive engineering and prototyping checks, with several subsystems moving to the pre-production stage. The first of the detector installations can be expected in late 2026 or 2027.

Chapter 4

Event Simulation and Reconstruction in CMS

Contents

4.1 Event Simulation in CMS 60
4.1.1 Hard Scattering and Parton Density Functions 60
4.1.2 Parton Showering and Hadronization 62
4.1.3 Underlying Event and Pileup 64
4.1.4 Event Generators 64
4.1.5 Detector Simulation 65
4.2 Reconstruction of Physics Objects 66
4.2.1 The Particle Flow Algorithm 67
4.2.2 Tracks 68
4.2.3 Primary Vertices 70
4.2.4 Muons 71
4.2.5 Electrons and photons 73
4.2.6 Jets 74
4.2.7 Tau Leptons 77
4.2.8 Missing Transverse Energy 80

To interpret the results of proton-proton collisions at the LHC, a detailed understanding of the underlying physical processes is crucial. Monte Carlo (MC) simulations provide a framework for modeling particle interactions based on theoretical predictions and simulating the detector's response. By comparing these simulated events with experimental data, physicists can validate the Standard Model (SM) or uncover deviations that may indicate new physics.

Reconstruction, on the other hand, serves as the vital bridge between raw detector data and meaningful physics insights. Using sophisticated algorithms, the electronic signals recorded by the CMS subdetectors are transformed into measurable quantities, such as particle energies and momenta.

This chapter explores the methodologies and tools used in the CMS experiment for simulating particle interactions and reconstructing events, forming the foundation for the analysis presented in subsequent chapters.

4.1 Event Simulation in CMS

Monte Carlo event generators build on the principle that physical systems can be described by probability density functions [122]. Using numerical MC methods to randomly sample these distributions, general-purpose MC event generators simulate particle collisions that resemble those observed in high-energy colliders like the LHC. The simulation of a proton-proton (pp) collision begins with the hard scattering of the partons (quarks and gluons) inside them. These high-momentum interactions result in several critical phenomena, including the production of the Higgs bosons and top quark pairs. At these high energy scales, the strength of the strong coupling constant (α_s) becomes sufficiently small, enabling the application of perturbative Quantum Chromodynamics (pQCD) to describe these processes effectively. However, the hard scatter is just the starting point of a more complex sequence of events. The scattered partons undergo further radiation of gluons and quarks, forming a cascade of secondary partons known as parton shower (PS). At lower energy scales, partons transition into color-neutral hadrons via hadronization, a phenomenon that is governed by non-perturbative QCD. Additionally, interactions unrelated to the primary hard scatter, collectively referred to as the underlying event (UE), contribute to the final-state particles. Monte Carlo event generators provide a detailed simulation of this entire collision sequence, from the initial hard scattering to the production of final-state particles. Figure 4.1 shows an illustration of a pp collision event simulated with an MC event generator.

Once the event generation is complete, detector simulations are used to model particle interactions with the detector components to replicate the signals they would produce in a real experimental environment.

The following subsections provide detailed information on MC event generation and detector simulation, addressing the individual stages of the collision process and the tools employed to simulate them.

4.1.1 Hard Scattering and Parton Density Functions

As stated, many processes of interest occurring at the LHC involve large momentum transfers. The perturbative nature of these interactions makes it possible to compute the hard scattering cross section $\hat{\sigma}$ using tools like Feynman diagrams.

While the hard scattering cross section is calculable perturbatively, it accounts only for the interaction between the incoming partons. The initial state of the hadrons and the distribution of their constituent partons cannot be determined by pQCD due to the non-perturbative nature of confinement in QCD. This information is instead encapsulated in parton density functions (PDFs), $f_a(x_a, \mu_F)$, which describe the probability of finding a parton a within a hadron, carrying a fraction x_a of the hadron's momentum at the factorization scale μ_F.

According to the factorization theorem [24,124], the total cross section for a process $\mathrm{p_1 p_2} \to n$ can then be expressed as a convolution of the non-perturbative PDFs and the perturbative

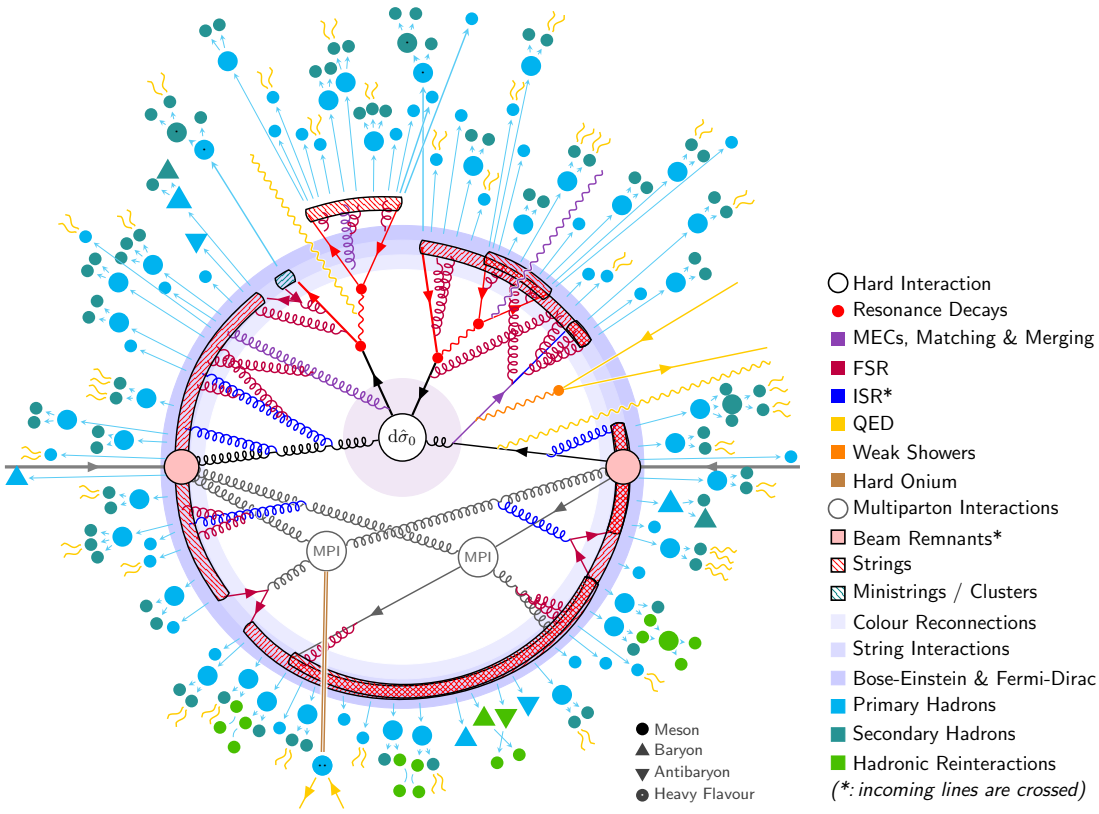

Figure 4.1: Sketch of a proton-proton collision as simulated by an MC event generator [123].

hard scattering cross section:

$$\sigma_{\mathrm{p_1p_2}\to n} = \sum_{a,b} \int_0^1 dx_a\, dx_b\, f_a^{\mathrm{p}_1}(x_a,\mu_F)\, f_b^{\mathrm{p}_2}(x_b,\mu_F)\, \hat{\sigma}_{ab\to n}(\mu_F,\mu_R) \tag{4.1}$$

$$= \sum_{a,b} \int_0^1 dx_a\, dx_b\, f_a^{\mathrm{p}_1}(x_a,\mu_F)\, f_b^{\mathrm{p}_2}(x_b,\mu_F)\, \frac{1}{2\hat{s}} \int d\Phi_n\, |\mathcal{M}_{ab\to n}|^2. \tag{4.2}$$

Here,

- x_a and x_b are the fractions of the proton momenta carried by the partons a and b. The indices a and b sum over all partons in the protons;
- μ_F and μ_R correspond to the factorization and renormalization scales, respectively;
- $f_a^{\mathrm{p}_1}(x_a,\mu_F)$ and $f_b^{\mathrm{p}_2}(x_b,\mu_F)$ are the PDFs for partons a and b, describing their momentum distributions inside the incoming protons p_1 and p_2, respectively;
- $\hat{\sigma}_{ab\to n}(\mu_F,\mu_R)$ is the parton-level cross section for the production of the final state n through the initial partons a and b. It depends on μ_F and μ_R and is related to the squared matrix element $|\mathcal{M}_{ab\to n}|^2$ as

$$\hat{\sigma}_{ab}(\hat{s}) = \frac{1}{2\hat{s}} \int d\Phi_n |\mathcal{M}_{ab\to X}|^2 \tag{4.3}$$

where $\hat{s}$ is the hadronic center-of-mass energy squared and $d\Phi_n$ represents the Lorentz-invariant phase space for the final state n.

MC generators calculate cross sections by incorporating PDFs to model the parton momentum distributions inside protons and computing matrix elements for the hard scattering process using pQCD. Matrix elements are calculated using Feynman rules, summing over all relevant diagrams, with techniques such as recursion relations and helicity amplitudes employed to manage computational complexity for multi-particle final states.

The factorization and renormalization scales, though theoretically arbitrary, are typically chosen to match the momentum transfer or center-of-mass energy of the process. To account for this arbitrariness, however, their values are varied up and down by a factor of 2 to estimate theoretical uncertainties.

The PDFs are generally extracted from global fits to experimental data, such as those provided by the H1 and ZEUS experiments at the HERA electron-proton collider and more recent measurements from CMS and ATLAS. To generate the simulations for the analysis detailed in this thesis, the NNPDF3.1 set [125] is used in the MC generators to account for parton density functions accurately. Figure 4.2 shows the NNPDF3.1 parton distribution functions at next-to-next-to-leading order (NNLO) accuracy for two different values of μ_F.

4.1.2 Parton Showering and Hadronization

Both the initial partons, as well as the quarks and gluons produced following a hard scattering process, can radiate additional partons through a mechanism known as parton showering. Unlike the calculable hard processes that involve large momentum transfers and are modeled using fixed-order perturbative QCD, parton showers involve multiple lower-energy emissions. Thus, MC generators employ dedicated PS algorithms to simulate the step-by-step evolution of each parton.

Central to parton shower simulation is the Sudakov form factor, which models the probability that a parton evolves from a higher scale t to a lower scale t' without undergoing splitting [126]. This form factor for a parton i can be mathematically expressed as:

$$\Delta(t, t') = \exp\left(- \int_{t'}^{t} \frac{dk}{k} \int_z dz \frac{\alpha_s(k)}{2\pi} P_{i\to jk}(z) \right) \tag{4.4}$$

where $P_{i\to jk}(z, k)$ are called splitting functions. These splitting functions dictate how momentum is distributed among the products of parton splitting.

In practice, within MC generators, the shower generation starts by randomly determining a lower scale t' for possible splitting using the Sudakov form factor. Once this factor suggests a high probability for splitting at scale t', a type of splitting, such as $q \to qg$ for quarks or $g \to gg$ and $g \to \mathrm{q\overline{q}}$ for gluons, is selected based on the relative probabilities provided by the splitting functions. Following the selection of the splitting type, the momentum fraction z is determined, dictating how the original parton's momentum is distributed among the newly

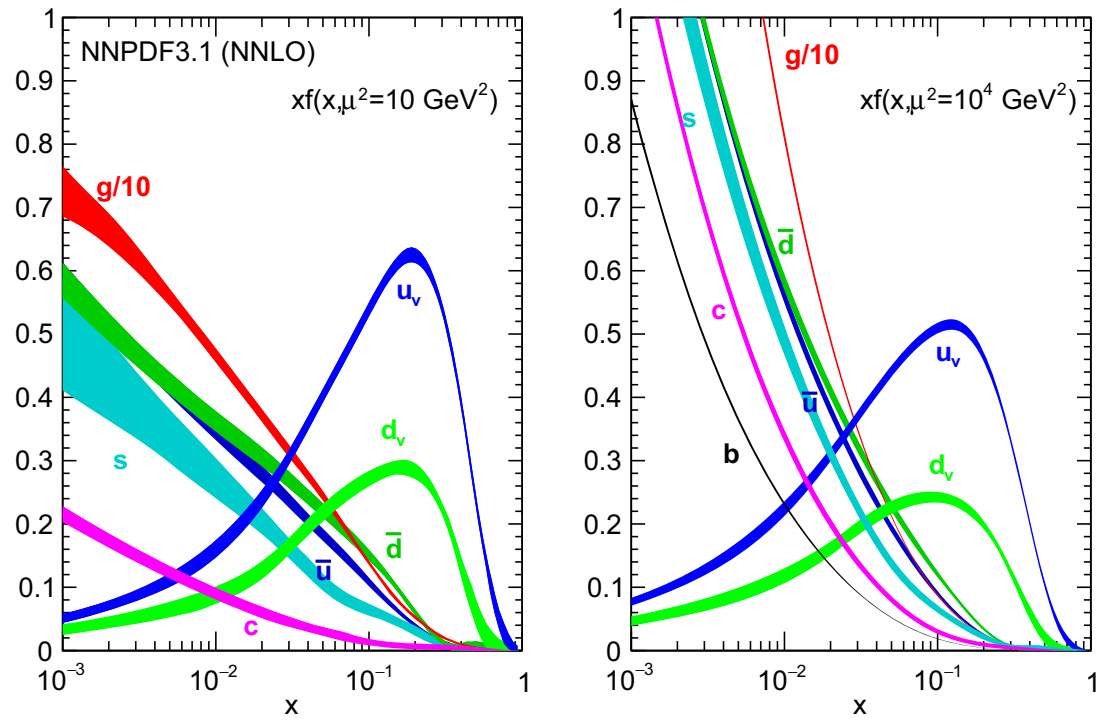

Figure 4.2: The NNPDF3.1 NNLO momentum-weighted parton distribution functions $xf(x, \mu^2)$ for various partons, shown as a function of the momentum fraction x carried by each parton type. The width of the curves depicts the uncertainties associated with the PDFs. The distributions are evaluated at two different factorization scales: $\mu_F^2 = 10\,\mathrm{GeV}^2$ (left) and $\mu_F^2 = 10^4\,\mathrm{GeV}^2$ (right). The gluon distribution (g) is scaled by a factor of 1/10 for better visualization [125].

created partons. This probabilistic process repeats for each new parton, continuously creating new branching points in the parton shower. The sequence repeats until the scale of parton energy falls below a certain threshold (around 1 GeV), at which point non-perturbative QCD effects take over.

Parton showers are categorized, based on when the splitting occurs, into either initial state radiation (ISR) or final state radiation (FSR). Initial state radiation is emitted by the partons involved in the hard interaction before the actual collision occurs and is typically modeled using a technique called *backward evolution*. This technique starts from the hard interaction and evolves backward towards the higher scales of the incoming particles. Conversely, FSR involves the radiation emitted by the products of the collision after the hard interaction and is modeled using *forward evolution*, wherein partons emerging from the collision evolve from the high-energy scales of the hard interaction down to lower energy scales.

As the energy of the partons decreases into the non-perturbative QCD regime, α_s becomes large, amplifying the effects of color confinement. Color confinement prevents the existence of free quarks and gluons. As a result, they are forced to combine into pairs and groups, forming new hadrons through a process known as hadronization.

Given that hadronization cannot be adequately described by pQCD methods, phenomenological models like the Lund string model [127,128] and the cluster model [129] are utilized. The Lund string model conceptualizes the space between a quark and an antiquark as a string that stretches as they move apart. When the energy of the partons drops below a certain threshold, the string becomes unstable and breaks, producing new quark-antiquark pairs that subsequently form hadrons. In the cluster model, partons first recombine into color-neutral clusters. These clusters either transform directly into hadrons or split into smaller clusters, which then decay into hadrons. The hadronization process ends once all partons have been converted into color-neutral hadrons.

4.1.3 Underlying Event and Pileup

It has been observed that hard scatter processes are accompanied by additional activity not attributed to ISR, FSR, or hadronization. This activity, collectively referred to as the underlying event, consists of multiple parton interactions (MPI), beam remnants, and soft radiation unassociated with the primary hard scatter. Monte Carlo generators simulate UE by calculating hard MPI using matrix element methods with parton density functions, while softer MPI is modeled through phenomenological approaches that incorporate the spatial overlap of colliding hadrons, often parameterized by the impact parameter. Beam remnants are modeled with energy-momentum conservation and quantum number assignment techniques. Further details on how UE is taken into account in MC generators can be found in Ref. [122]. To describe the underlying events in the simulated samples of the $H \to a_1 a_1$ analysis, the CP5 tune [130] derived from the combined fits to the CDF UE data at $\sqrt{s} = 1.96\,\text{TeV}$ and CMS UE data at $\sqrt{s} = 7\,\text{TeV}$ and $13\,\text{TeV}$ is used.

While underlying events describe the secondary processes in a single hard scattering event, pileup refers to additional collisions occurring within the same or adjacent bunch crossings in a collider experiment. Pileup can introduce extraneous signals that can interfere with the detector response to the hard scatter products. The simulation of pileup in the CMS experiment begins by deriving the average number of pileup interactions per bunch crossing (N) from LHC collision data, as shown in Fig. 3.4. Minimum bias events, representing typical inelastic collisions, and primary hard scattering events are simulated separately. A Poisson distribution with mean N is then sampled to determine the number of minimum bias events to overlay on each hard scatter. These events are combined at the detector simulation level to produce a dataset that accurately incorporates pileup effects.

4.1.4 Event Generators

A number of MC event generators, each focusing on one or a few of the aspects of the collision process described above, are used in CMS to simulate a wide range of SM and beyond-SM events. The simulated events used for the physics analysis presented in this thesis have been produced with the following generators:

- PYTHIA (v.8.212) [131,132] is a general-purpose generator designed to simulate every stage of the collision process, from the initial hard scattering to the final hadronization. It performs matrix element calculations for the hard subprocess at leading-order (LO) accuracy and is particularly notable for its advanced parton shower and hadronization techniques.

PYTHIA employs a p_T-ordered dipole showering approach to model both initial and final state radiation effectively. It also includes a detailed simulation of multiple parton interactions and uses the Lund string model for hadronization.

- MADGRAPH5_aMC@NLO (v.2.6.5) [133] is an event generator specialized in the automated computation of matrix elements for $2 \to n$ processes at both leading order (LO) and next-to-leading order (NLO) accuracy. While it excels in generating the hard scattering process, MADGRAPH5_aMC@NLO does not simulate parton showers or hadronization and is typically interfaced with PYTHIA for a complete simulation of particle collisions. To ensure proper integration, jet matching and merging techniques are employed to prevent the double-counting of jets or gaps in the phase space. For LO simulations, the MLM scheme is used, where events with partons exceeding the matrix element's p_T coverage are discarded [134]. For NLO simulations, the FxFx method is applied, which accounts for overlapping contributions from identical diagrams at different orders [135].
- POWHEG BOX (v.2.0) [136–138] is an advanced event generator designed for precise NLO matrix element calculations. It implements the Positive Weight Hardest Emission Generator (POWHEG) method to perform the matching of matrix elements with parton showers.

4.1.5 Detector Simulation

While MC generators provide detailed kinematic information for final-state particles, a realistic picture requires passing these events through a detailed detector simulation to model their interactions with the detector accurately. For the datasets used in this thesis—and as standard practice within the CMS experiment—the GEANT4 simulation package is employed [139, 140]. GEANT4 meticulously models the complex geometry of the CMS detector, incorporating its active, passive, and inert materials (cables, cooling pipes, and structural components) while also simulating the propagation of charged particles in the electric and magnetic fields. It performs detailed tracking of each charged particle's trajectory, recording detector hits and interaction vertices with precision. The package calculates energy deposits in the calorimeters, accounting for the interactions of electrons, photons, and hadrons with absorber materials, and simulates the resulting electronic signals. It also handles the conversion of analog signals to digital outputs, ensuring an accurate representation of the detector's readout and electronic response. To maintain fidelity, the detector geometry in GEANT4 is updated for each data-taking period to reflect the actual conditions within the detector at the time. For scenarios where speed and computational efficiency are critical, the CMS experiment employs Fast Simulation (FASTSIM), which simplifies the modeling of particle interactions with the detector by using parameterized responses [141–144].

With the detector simulation stage, the event simulation process is complete. The outputs, whether generated by GEANT4 or FASTSIM, are in the same format and contain equivalent information as those recorded by the CMS detector during data-taking. This enables the use of the same reconstruction and identification algorithms for both simulated events and real data to reconstruct physics objects.

4.2 Reconstruction of Physics Objects

The CMS detector records a diverse array of raw readout data as particles traverse and interact with its various subsystems. Figure 4.3 illustrates a transverse slice of the CMS detector, showing how different particles leave characteristic signatures as they pass through the silicon tracker, electromagnetic and hadronic calorimeters, and the muon chambers. These raw signals are then reconstructed into distinct particles using the Particle Flow (PF) algorithm by integrating inputs from all CMS subsystems [123]. A comprehensive overview of this algorithm is provided in Sec. 4.2.1.

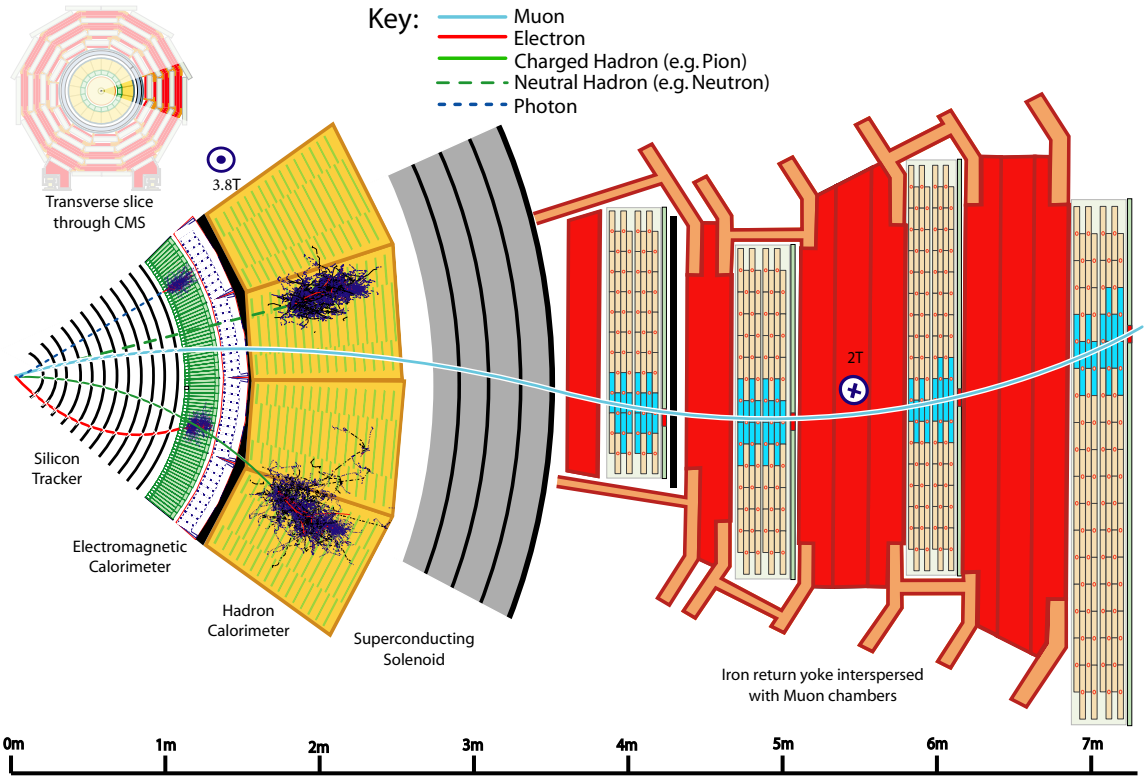

Figure 4.3: A sketch of the interactions of different particles in a transverse slice of the CMS detector, as they traverse from the beam interaction region to the muon detector [123].

The reconstruction process begins with the generation of fundamental components—tracks and energy clusters. Tracks, which map the trajectories of charged particles, are reconstructed using data from the CMS tracker and muon system. Simultaneously, energy clusters resulting from the interaction of electrons, photons, and hadrons are identified within the Electromagnetic Calorimeter (ECAL) and the Hadron Calorimeter (HCAL). These tracks and clusters are carefully combined by the PF algorithm to accurately reconstruct identifiable particles such as muons, electrons, photons, and hadrons. The data is further processed to assemble higher-level objects such as tau leptons and jets. Additionally, calculations of missing transverse energy ($\vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}$) are performed, which helps to identify particles, like neutrinos, that interact weakly with the detector. The methodologies for reconstructing each of these so-called "physics objects" are explained in detail in Sections 4.2.2-4.2.8.

4.2.1 The Particle Flow Algorithm

The key advantage of the PF algorithm lies in its ability to combine information from various detector components, which helps achieve superior reconstruction and identification performance compared to methods relying on a single subdetector.

The PF reconstruction starts by identifying low-level detector signals, known as PF elements, which include tracks in the tracker and clusters in the calorimeters. Clustering in the CMS calorimeters is carried out independently in the ECAL, HCAL, and preshower layers using a dedicated algorithm [123]. The process begins by identifying seed cells based on a predefined energy threshold and then forming topological clusters by grouping these seeds with neighboring cells that have energy exceeding twice the expected noise level. Once the topological clusters are formed, an expectation-maximization algorithm based on a Gaussian mixture model is applied. It models the energy measured in each calorimeter cell as a sum of contributions from several overlapping Gaussian functions, each associated with a nearby seed location. This algorithm iteratively refines the positions and energies of these functions until their parameters accurately reflect the observed energy patterns in the calorimeters. The reconstruction of tracks, on the other hand, proceeds as detailed in Sec. 4.2.2.

The next step involves linking these PF elements using their spatial and kinematic relationships to form PF *blocks*, which represent groups of interconnected detector signals likely associated with the same particle. In this process, tracks are extrapolated to calorimeters to link with clusters, clusters across calorimeters are matched based on spatial overlap, and muon tracks are associated with corresponding signals in the muon system.

Once the PF blocks are constructed, the algorithm proceeds with particle identification in a structured sequence:

- **Muons** are first identified by matching tracks from the silicon tracker with tracks in the muon system and calorimeter clusters, compatible with the signal from minimum ionizing particles.
- **Electrons** are reconstructed by linking tracks from the tracker to energy clusters in the ECAL.
- **Photons** are identified as ECAL clusters that are not linked to any tracks.
- **Charged hadrons** are reconstructed using the remaining tracks associated with energy deposits in both the ECAL and HCAL.
- Lastly, **neutral hadrons** are identified as the remaining HCAL clusters that are not linked to any tracks.

As each particle type is reconstructed, its associated tracks and clusters are removed from further consideration. The reconstructed particles are subsequently used to build complex physics objects such as jets and hadronically decaying taus, as well as to measure $\vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}$. Detailed discussions of how these specific physics objects are reconstructed will follow in the subsequent subsections.

4.2.2 Tracks

Track reconstruction, or tracking, is performed at the CMS experiment using the CMS tracking software commonly referred to as the Combinatorial Track Finder (CTF) [145]. The CTF, based on combinatorial Kalman filtering [146–150], builds tracks in four steps:

- The **seed generation** step involves generating initial track candidates, or seeds, which define the preliminary trajectory parameters and their uncertainties. Seeds are formed based on two main parameters: the seeding layers, which are specific detector layers where hits are searched for, and the tracking regions, which define acceptable track parameters, such as minimum p_{T} or proximity to the beam spot[1]. The seeding layers are usually pairs, triplets, or quadruplets of detector layers. Seed formation typically occurs in the pixel detector due to its high granularity and low occupancy, which provide precise estimates of trajectory parameters and facilitate efficient track reconstruction. Additionally, seeds are also created in the strip detector to identify tracks originating outside the pixel volume.

- In the **track finding** step, seed trajectories are extrapolated using a Kalman filter (KF) along a charged particle's expected flight path. The process uses initial estimates from the seed to search for additional hits in adjacent detector layers. If the found hits pass a χ^2 compatibility test, the algorithm includes them in the track and updates the respective trajectory parameters. If no hit is found due to module inefficiency, a "ghost hit" may be added to continue the search. This step continues until no more compatible layers are found, multiple missing hits occur, or the track's p_{T} falls below a set threshold.

- After collecting all the hits, the **track fitting** step employs a Kalman Filter (KF) to refine the initial estimates of track parameters. The KF begins by processing the innermost hits. The fit then advances iteratively to the outermost hits, with the position and uncertainty of each valid hit recalculated based on the updated track parameters. The trajectory is further refined using a smoothing technique where a second filter runs in reverse from the final hit toward the beam line. This bidirectional approach ensures precise parameter estimation at all points, with optimal parameters for each track surface derived from the weighted average of the two filters.

- The final step involves **track selection**, where tracks are classified into *loose*, *tight*, or *high-purity* categories. The loose criteria define the minimum requirements for a track to be retained, while the tight and high-purity flags provide progressively stricter requirements to improve track purity further. The track selection procedure was gradually improved over time, starting with a parametric selection in LHC Run 1 and then moving to a boosted decision tree (BDT) in Run 2.

The final collection of reconstructed tracks is produced after multiple iterations of the CTF track reconstruction sequence. This process, called *iterative tracking*, starts with searching for tracks that are easiest to find, like those with relatively large p_{T}, and produced near the interaction region. With each iteration, the hits linked to the identified tracks are removed to reduce combinatorial complexity and enhance the detection of harder-to-find tracks, such as those with significant displacement. Over time, these iterations have been constantly

[1]The beam spot refers to the three-dimensional luminous region within a particle detector where the proton beams from a bunch crossing intersect.

developed to adapt to diverse needs, including adjustments made to accommodate the new pixel layer added during the Phase 1 tracker upgrade.

The performance of the tracking algorithm is primarily evaluated based on tracking efficiency. The efficiency achieved during Run 2 is highlighted in Figures 4.4 and 4.5. Specifically, Fig. 4.4 illustrates the efficiency for non-isolated charged particles from a simulated $t\bar{t}$ sample under various pileup conditions [151]. At an average pileup of 35, the efficiency peaks at 94% within the tracker barrel ($|\eta| < 0.9\,\text{GeV}$) and varies between 84-88% at higher pseudorapidities, primarily due to the nuclear interactions of hadrons within the tracker material. In contrast, Fig. 4.5 presents the efficiency for isolated muons, measured using the tag-and-probe method from $Z \to \mu\mu$ events and Drell-Yan MC simulations [152]. Here, the efficiency is nearly 100% across the entire detector, attributable to the straightforward interactions of muons with the silicon detector through ionization and their minimal energy loss from bremsstrahlung.

During LS2, several developments were undertaken to improve the tracking algorithm, mainly focusing on the reduction of reconstruction time and fake rate. A new algorithm, mkFit, was introduced for track pattern recognition that exploits parallelization and vectorization in multi-core CPU architectures [153]. This algorithm is now implemented in a subset of the CMS tracking iterations and maintains comparable physics performance to the traditional CKF-based approach while improving computational efficiency. Additionally, for Run 3, the track selection procedure was upgraded from the BDT-based method to a Deep Neural Network (DNN)-based approach, further improving the selection process [154].

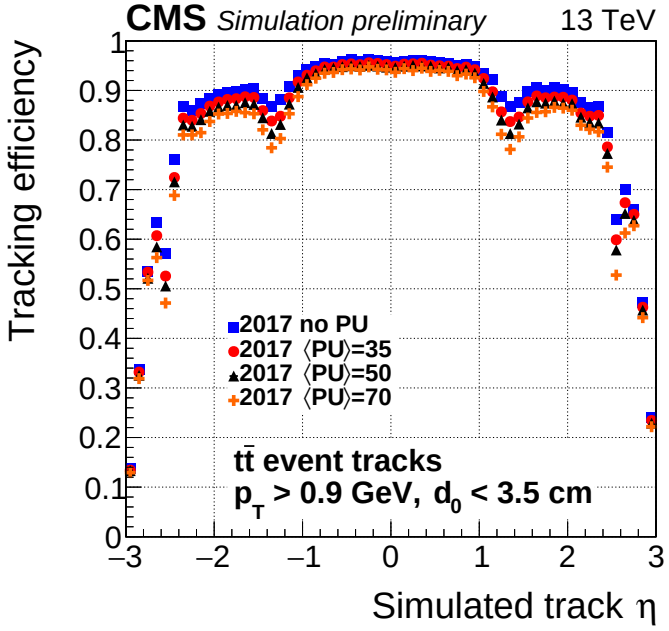

Figure 4.4: Track reconstruction efficiency as a function of track η, at different pileup conditions, for a sample of simulated $t\bar{t}$ events. The efficiency is shown for tracks that meet the high-purity criteria and have $p_T > 0.9\,\text{GeV}$ [151].

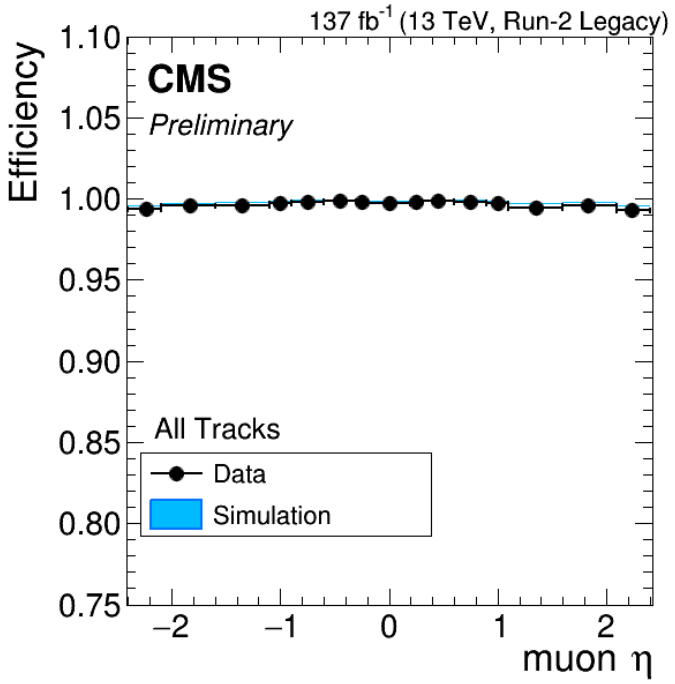

Figure 4.5: Tracking efficiency measured from $Z \to \mu\mu$ events with the tag-and-probe method, as a function of η of the probe muon [152].

Tracks play a crucial role in the $H \to a_1 a_1$ analysis central to this thesis, as detailed in Chapter 5. The high-purity track collection is specifically utilized to ensure that only data of the highest quality and reliability are used.

4.2.3 Primary Vertices

At the LHC, multiple proton-proton interactions can occur during a single bunch crossing, leading to several vertices within an event. The primary vertex (PV) reconstruction at CMS is designed to identify and accurately determine the positions of these interaction vertices, along with their associated uncertainties [145].

Before proceeding with the description of the vertex reconstruction procedure, it is important to note that this process relies on precise knowledge of the beam spot. The beam spot serves as a critical reference for clustering tracks and identifying interaction vertices. The center of the beam spot (x_{BS}, y_{BS}, z_{BS}) is determined either by averaging the positions of reconstructed primary vertices across multiple events or through a χ^2-fit that exploits correlations between the transverse impact parameters (d_0) of tracks and their azimuthal angles (ϕ):

$$d_0(\phi, z_0) = x_{BS} \sin\phi + \frac{dx}{dz} \sin\phi(z_0 - z_{BS}) - y_{BS} \cos\phi - \frac{dy}{dz} \cos\phi(z_0 - z_{BS}), \tag{4.5}$$

where $\frac{dx}{dz}$ and $\frac{dy}{dz}$ represent the slopes of x and y relative to z, respectively, and z_0 corresponds to the longitudinal impact parameter of the tracks.

The transverse (σ_x, σ_y) and longitudinal (σ_z) spreads of the beam spot are estimated using

a likelihood fit of the primary vertex distribution or via correlations between the transverse impact parameters of track pairs originating from the same vertex:

$$\left\langle d_0^{(1)} d_0^{(2)} \right\rangle = \frac{\sigma_x^2 + \sigma_y^2}{2} \cos(\phi_1 - \phi_2) + \frac{\sigma_y^2 - \sigma_x^2}{2} \cos(\phi_1 + \phi_2). \tag{4.6}$$

These parameters are updated every luminosity section (23 seconds) to account for potential drifts or variations in the beam position and size.

With this foundation in place, the PV reconstruction process proceeds in three main steps:

1. **Track Selection:** Tracks consistent with originating from the primary interaction region are selected based on criteria such as the transverse impact parameter significance with respect to the beam spot, the number of pixel and strip hits associated with the track, and the normalized χ^2 of the track fit.

2. **Track Clustering:** The selected tracks are grouped into clusters based on their z-coordinates at their points of closest approach to the beam spot using the Deterministic Annealing (DA) algorithm [155].

3. **Vertex Fitting:** Vertex candidates with at least two associated tracks are fitted using an adaptive vertex fitter [156], which computes precise vertex parameters, including the x, y, and z positions and the covariance matrix. Each track is assigned a weight based on its compatibility with the vertex, with weights close to 1 indicating strong compatibility.

Reconstructed vertices within an event are then ranked by the quadratic sum of the transverse momenta, $\sum p_{\mathrm{T}}^2$, computed using jets (formed by clustering tracks with the anti-k_{T} algorithm [157]), remaining single tracks (including leptons), and $\vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}$. The primary vertex with the highest $\sum p_{\mathrm{T}}^2$ value is identified as the hard-scatter vertex, while the rest are classified as pileup vertices [89].

4.2.4 Muons

The CMS experiment is designed to achieve exceptional muon reconstruction performance. Both the silicon tracker and muon chambers play a central role in this process. The muon reconstruction algorithm begins with a local reconstruction step, where individual hits in the DTs and CSCs are combined into three-dimensional segments [111, 158]. These segments are then used as seeds for a Kalman filter, which reconstructs the muon trajectory by incorporating all available hits from the DT, RPC, and CSC detectors. The resulting collection of muon tracks is referred to as *standalone muons*. While standalone muons are rarely used in physics analyses due to significant contamination from cosmic ray backgrounds, they can be matched with tracks reconstructed in the silicon tracker to create two other distinct classes of muon tracks:

- Global muon: Constructed using an outside-in approach, these tracks begin as standalone muon tracks, which are matched to inner tracker tracks by evaluating the compatibility of their parameters when extrapolated to a common surface. The hits from the standalone and tracker tracks are then combined, and a final fit is performed to produce the global muon track.

- Tracker muon: These tracks are reconstructed using an inside-out approach by extrapolating tracks from the inner tracker with $p_{\mathrm{T}} > 0.5\,\mathrm{GeV}$ and total momentum $p > 2.5\,\mathrm{GeV}$ to the muon system. A tracker track qualifies as a tracker muon if it matches with any DT or CSC segment, i.e. if the absolute difference in their x-coordinate is less than 3 cm or the distance-to-uncertainty ratio is below 4.

As both the tracker and muon systems guarantee high reconstruction efficiency, approximately 99% of muons produced within the geometrical acceptance of the muon system are reconstructed as either global-muon tracks or tracker-muon tracks, and very often as both. Global and tracker muons sharing the same tracker track are merged into a single candidate.

Additional identification (ID) criteria are applied to muon candidates to classify them based on their reconstruction quality. These criteria rely on variables associated with the muon reconstruction, such as the χ^2 of the track fit and the number of hits in the inner tracker and the muon system. A key factor is the muon segment compatibility, which measures how well an extrapolated tracker track matches muon segments and provides scores ranging from 0 to 1. Additional checks include a kink-finding algorithm, which splits the tracker track at several points and compares the resulting segments to determine if they are consistent with a single track. External factors, such as the muon's compatibility with the primary vertex, are also used in the selection process. Using these criteria, CMS has defined the following standard muon ID working points (WPs):

- Loose muon ID: Loose muons are muons selected by the PF algorithm that are also either tracker or global muons. This ID primarily targets prompt muons originating at the primary vertex and muons from light or heavy-flavor hadron decays.

- Medium muon ID: A medium muon must satisfy the Loose ID criteria, with the additional requirement that their tracker track uses hits from more than 80% of the traversed inner tracker layers. Further selection criteria depend on whether the muon is reconstructed as only a tracker muon or as both a tracker and global muon, as summarized in Table 4.1. This ID is optimized for prompt muons and muons from heavy-flavor decays.

- Tight muon ID: Tight muons must meet stricter requirements, including tracker tracks with hits from at least six tracker layers, segment matching in two muon stations, and compatibility with the primary vertex ($|d_0| < 0.2$ cm, $|d_z| < 0.5$ cm). The muons must be reconstructed as both tracker and global muons. This ID is designed to suppress muons from decay-in-flight and hadronic punch-through.

Muon type	Criteria	Requirement
Tracker muon	Muon segment compatibility	> 0.451
Tracker & Global muon	Muon segment compatibility	> 0.303
	Global fit χ^2/dof	< 3
	Position match χ^2 (tracker and standalone muon)	< 12
	Kink-finding algorithm χ^2	< 20

Table 4.1: Selection criteria for the medium muon ID, in addition to the requirement that the tracker track of the muon uses hits from more than 80% of the traversed inner tracker layers. These criteria are categorized based on whether the muon is reconstructed as only a tracker muon or as both tracker and global muons.

In addition to these three muon IDs, CMS employs a variety of specialized IDs optimized for cases such as high-pt muons ($p_T > 200\,\text{GeV}$) or low-pt muons coming from B-hadron decays. More recently, MVA-based identification techniques have been introduced, as described in Ref. [159], to enhance signal efficiency and background rejection.

For selecting muons in the $H \to a_1 a_1$ analysis described in Chapter 5, the Medium WP was utilized. The efficiency of this WP with respect to the muon p_T and η is shown in Fig. 4.6, as calculated using the tag-and-probe method with data and simulation from 2017 [160]. The efficiency exceeds 98% across the p_T range 20-200 GeV, and the data and simulation agree to within 1%. The efficiency varies between 96% and 99%, depending on η, and the data and simulation agree to within 1-3%. The efficiency drop around $|\eta| = 0.2$ is attributed to cracks between wheels in the muon detector, while the dip in the forward region arises from inactive muon chambers not being modeled in the MC simulation.

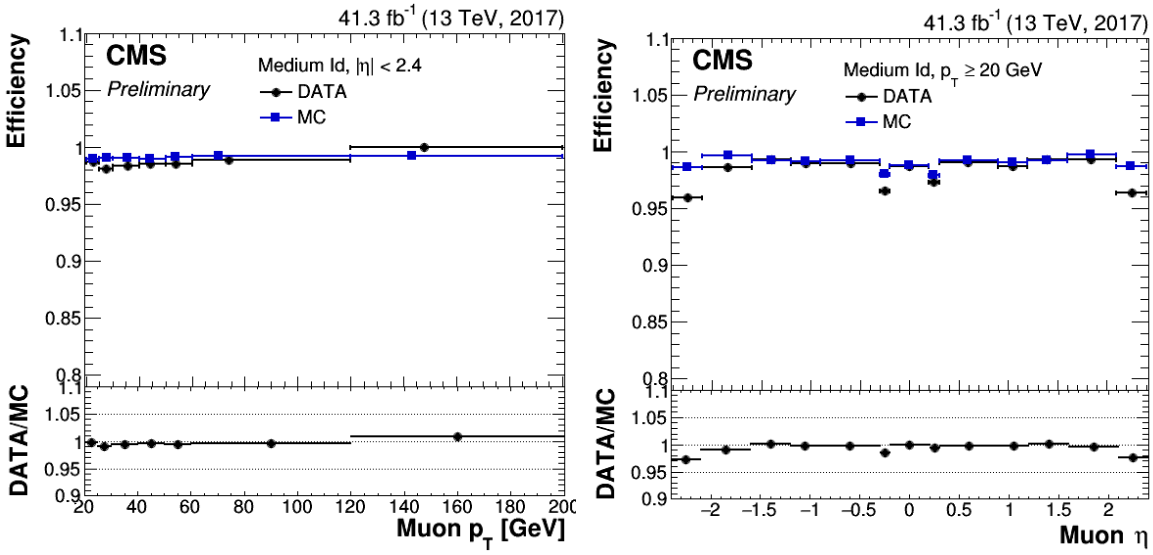

Figure 4.6: Muon identification efficiency for the Medium working point as a function of the muon transverse momentum (left) and pseudorapidity (right). The efficiency is studied with the tag-and-probe method applied to data and simulated $Z \to \mu\mu$ events for the year 2017 [160].

4.2.5 Electrons and photons

Electrons can interact extensively with the tracker material and lose energy through the emission of bremsstrahlung photons. Similarly, high-energy photons may convert into electron-positron pairs, which subsequently emit bremsstrahlung photons, creating a cascading effect. Therefore, by the time electrons and photons reach the ECAL, they may have evolved into a shower of secondary particles. These interactions make electron and photon reconstruction particularly challenging, thereby requiring specialized algorithms for the task.

The reconstruction of electrons and photons in CMS begins with the formation of clusters by grouping together crystals with energies exceeding $\sim 80\,\text{MeV}$ in the ECAL barrel and $\sim 300\,\text{MeV}$ in the endcaps [161]. The cluster with the highest energy deposited in a certain η-ϕ region is identified as a seed cluster, and the surrounding clusters within a defined geometric

window are then merged with it to form superclusters (SC).

Additionally, tracks for electrons are reconstructed using two different seeding techniques [162]. The ECAL-driven seeding, optimized for isolated electrons with high transverse energy, extrapolates a trajectory from the SC toward the primary vertex and accepts tracks with hits matching the predicted path. The tracker-driven method, designed for low-p_{T} electrons, matches loosely selected SCs with inner tracker tracks using cut-based and multivariate criteria. Once seeded, the electron trajectory is iteratively built using a Kalman Filter, with energy loss modeled by a Bethe-Heitler distribution. For the final track fit, a Gaussian Sum Filter (GSF) is used instead of the traditional CTF, which accounts for the significant bremsstrahlung energy losses by approximating energy loss as a superposition of Gaussian distributions [162, 163]. The GSF tracks thus obtained are then extrapolated toward the ECAL for track-cluster association.

The ECAL superclusters and GSF tracks are linked in the PF algorithm to form particle blocks, which are then resolved into e/γ objects: SCs with GSF tracks are identified as electrons, while those without are identified as photons. Additional selection criteria are applied to these objects to ensure accurate identification and isolation, as explained in Ref. [161].

4.2.6 Jets

Due to color confinement, quarks and gluons produced in collisions cannot exist freely; instead, they fragment and hadronize into collimated particle showers, known as jets. Jets are an inevitable feature of LHC collisions, making their accurate reconstruction essential for studying the properties of the originating partons.

The standard jet reconstruction in CMS is performed by clustering together PF candidates using the anti-k_{T} algorithm [157] implemented in the FastJet software package [164]. The algorithm defines two distance parameters:

- Distance between object i and object j, d_{ij}:

$$d_{ij} = \min(p_{\mathrm{T},i}^{-2}, p_{\mathrm{T},j}^{-2}) \frac{\Delta R_{ij}^2}{R^2} \tag{4.7}$$

  Here, $\Delta R_{ij} = \sqrt{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}$ is the distance between the objects i and j in the y-ϕ space. The parameters $p_{\mathrm{T},i}$, y_i and ϕ_i represent the transverse momentum, rapidity and azimuth of particle i. R is the jet radius parameter that defines the maximum size of the reconstructed jets. In CMS, $R = 0.5$ was used as the standard during Run 1 [165], and $R = 0.4$ has been the default choice since Run 2 [166]. Larger values, such as $R = 0.8$, are employed for boosted topologies, like hadronic W-boson decays, which result in jets with overlapping substructures.

- Distance between object i and the beam B, d_{iB}:

$$d_{iB} = p_{\mathrm{T},i}^{-2} \tag{4.8}$$

The algorithm then proceeds iteratively by identifying the smallest distance d_{min} among all d_{ij} and d_{iB}:

- If $d_{\text{min}} = d_{ij}$, objects i and j are merged into a single pseudojet by summing their four-momenta.
- If $d_{\text{min}} = d_{iB}$, object i is declared a final jet and removed from the list.

This process is repeated until all PF objects have been assigned to jets. The final jet momentum is accordingly defined as the sum of its constituents' momentum.

Notably, the anti-k_{T} algorithm is both *infrared-safe*, ensuring stability in the presence of very low-energy particles, and *collinear-safe*, meaning it remains unaffected by particles splitting along their flight path. These features make the algorithm particularly resilient to soft radiation and pileup noise, establishing it as the preferred technique in many experiments.

To reject poorly reconstructed jets and jets arising from instrumental noise and detector inefficiencies, the clustered jets are required to meet some identification criteria as described in Refs. [165,166]. These criteria rely on the charged and neutral particle multiplicities within the jet, as well as the fractions of jet energy contributed by various PF particle types. For the analysis described in this thesis, the *tight* PF jet ID working point is used, which provides a noise rejection rate exceeding 99.9% in the central η region [166].

Pileup mitigation for jets

Pileup poses a significant challenge in jet reconstruction at the LHC, introducing additional particles that can contaminate jets and distort measurements of their energy and momentum. To address this, CMS employs two primary pileup mitigation algorithms—Charged Hadron Subtraction (CHS) and PileUp Per Particle Identification (PUPPI)—both of which are applied to PF candidates before jet clustering [123,167].

The CHS method utilizes tracking information to identify and remove charged-particle candidates associated with pileup vertices. Only charged hadrons linked to the primary vertex are retained for jet clustering, effectively reducing pileup contamination. However, CHS is limited to the tracker region and does not address contributions from neutral pileup particles.

In contrast, PUPPI operates by assigning weights to PF candidates based on their likelihood of originating from pileup. These weights are determined using local and global event properties, such as the particle's proximity to the primary vertex, p_{T}, and the local pileup density. Particles likely to originate from pileup interactions are assigned lower weights or excluded entirely from jet reconstruction. Unlike CHS, PUPPI accounts for both charged and neutral pileup contributions, providing superior suppression in high-density environments.

In Run 2, CMS adopted the CHS method for pileup mitigation in narrow jets and PUPPI for large-area jets. From Run 3 onward, PUPPI has become the default algorithm for both narrow and large-area jets. For the $\text{H} \to \text{a}_1\text{a}_1$ analysis data described in this thesis, performed using Run 2 data, the CHS technique is employed.

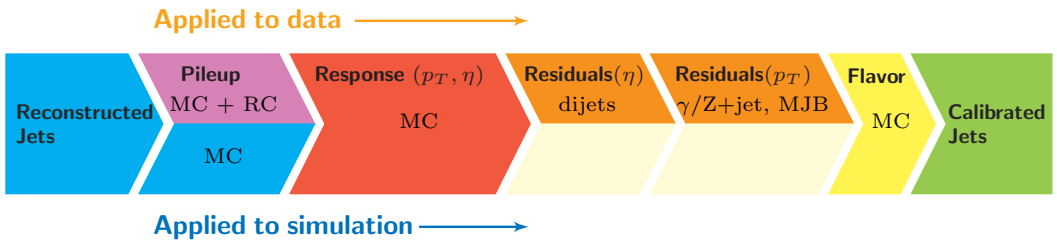

Figure 4.7: Illustration of the jet energy calibration procedure used in data and simulation [168].

Jet Energy Corrections

The energy of jets measured in the detector often deviates from the corresponding particle jet energy obtained from simulation using the same jet algorithm. This mismatch arises due to the detector's non-uniform and non-linear response, as well as additional contributions from electronic noise and pileup. Jet energy corrections are, therefore, crucial to calibrate the reconstructed jet energies to match the actual particle jet energies.

CMS uses a multi-step, factorized approach to calibrate jet energies, as illustrated in Fig. 4.7 [165, 168]. The procedure includes:

- Pileup Offset Corrections: This initial step removes residual contributions from pileup, particularly from neutral particles. The offset is estimated by comparing jet p_{T} in simulated QCD dijet events with and without pileup overlay as a function of p_{T} and η. Residual corrections are derived for data using the random cone method, which averages the momenta of PF candidates inside randomly placed cones in zero bias events and simulated samples.
- Response Corrections: These corrections address the detector's non-uniform and non-linear energy response by comparing reconstructed jet energy to generator-level jet energy in simulated QCD multijet events.
- Residual Corrections for data: These corrections account for any small discrepancies between data and simulation and are applied in two steps:
  - Relative η-dependent corrections, which normalize the jet response across different η regions relative to the barrel region.
  - Absolute p_{T}-dependent corrections, derived using momentum-balance techniques to cover the full p_{T} range.
- Flavor-Dependent Corrections (Optional): These corrections adjust for differences in the detector's response to quark-initiated and gluon-initiated jets.

The performance of the jet energy corrections derived with 13 TeV data collected during LHC Run 2 can be found in Ref. [169].

Heavy Flavor Tagging

Jets originating from the hadronization of b- and c-quarks are an important signature in many physics analyses, making their identification highly significant. Heavy-flavor hadrons present in these jets are characterized by relatively long lifetimes, of the order of $10^{-13} - 10^{-12}$ seconds [20]. Depending on their momentum, b-hadrons can travel up to a few centimeters and c-hadrons several millimeters before decaying, producing displaced tracks and forming a secondary vertex, as illustrated in Fig. 4.8. Flavor-tagging algorithms exploit these distinctive features to calculate the likelihood that a jet originates from a heavy-flavor quark.

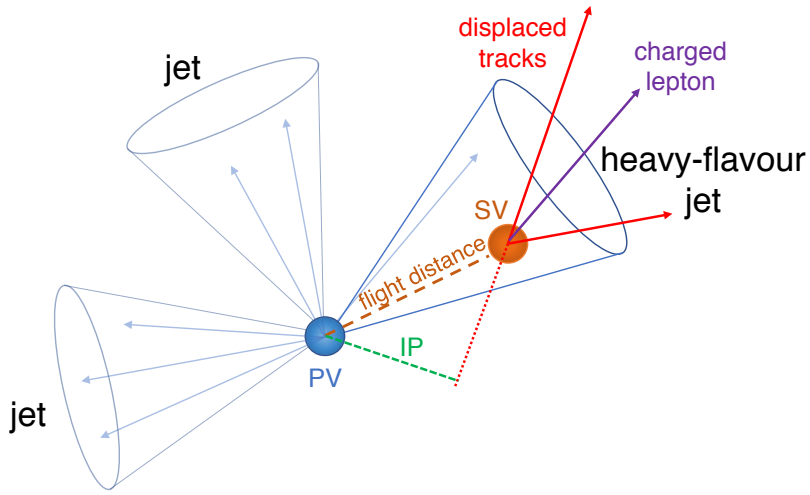

Figure 4.8: Illustration of a heavy-flavor jet showing a secondary vertex (SV) from the decay of a b- or c-hadron, producing displaced charged-particle tracks [170].

In the analysis presented in this thesis, which focuses on the $H \to a_1 a_1$ decay with two same-sign muons in the final state, QCD multijet events with heavy hadrons that decay into muons constitute a major background. To reduce the contribution from these events, a veto on b jets is applied. These b jets are identified using the DeepJet algorithm [171, 172]. DeepJet employs advanced machine-learning techniques based on DNNs and takes a wide range of features as input, including properties of the individual jet constituents, displaced tracks, secondary vertices, and global jet observables. It outputs probabilities for each jet flavor (b-, c-, light-flavor, or gluon) that are combined into a b tagging discriminator. CMS defines standard working points—loose, medium, and tight—based on thresholds applied to this discriminator. For the purposes of this analysis, the tight working point is used. This WP corresponds to a misidentification rate of 0.1% (1%) for light-flavor quarks or gluons (c-jets) [170] and a b tagging efficiency of approximately 65% measured in $t\bar{t}$+jets events [173]. Events containing one or more b-tagged jets with $p_T > 20\,\text{GeV}$ and $|\eta| < 2.4$ are vetoed.

4.2.7 Tau Leptons

Tau leptons are unstable particles with a measured lifetime of $(290.3 \pm 0.5) \times 10^{-15}$ s [20]. When produced during high-energy collisions at the LHC, they undergo rapid decays before reaching the detectors, thus necessitating their indirect identification through decay products. As the

Decay mode	Resonance	$\mathcal{B}$ (%)
Leptonic decays		
$\tau^- \to e^- \bar{\nu}_e \nu_\tau$		17.8
$\tau^- \to \mu^- \bar{\nu}_\mu \nu_\tau$		17.4
Hadronic decays		
$\tau^- \to h^- \nu_\tau$		11.5
$\tau^- \to h^- \pi^0 \nu_\tau$	$\rho(770)$	25.9
$\tau^- \to h^- \pi^0 \pi^0 \nu_\tau$	$a_1(1260)$	9.5
$\tau^- \to h^- h^+ h^- \nu_\tau$	$a_1(1260)$	9.8
$\tau^- \to h^- h^+ h^- \pi^0 \nu_\tau$		4.8
Other modes with hadrons		3.3

Table 4.2: Different decay modes of the tau lepton with their corresponding branching fractions ($\mathcal{B}$) in % [20]. The known intermediate hadron resonances are listed where relevant. Charged hadrons are denoted by the symbol $h^\pm$. In about 98% of cases, $h^\pm$ are charged pions $\pi^\pm$, and in the remaining 2% kaons, $K^\pm$. Although only τ^- decays are presented, the corresponding charge-conjugate decays share identical decay channels and branching fractions.

heaviest lepton in the SM, the tau lepton can decay into a variety of lighter particles, either leptonically, into electrons or muons accompanied by neutrinos or hadronically into hadrons (sometimes via intermediate mesonic resonances) along with a tau neutrino. Leptonic tau decays account for approximately 35% of the cases, while hadronic tau decays make up the remaining 65%, typically yielding either one or three charged hadrons (mostly pions, $\pi^\pm$, and kaons, $K^\pm$) and up to two neutral pions (π^0). These decay modes and their branching fractions are detailed in Tab. 4.2.

For electrons and muons originating from τ decays, reconstruction is performed using the standard PF approach as detailed in Sections 4.2.5 and 4.2.4. Distinguishing these leptons from prompt leptons presents a significant challenge, hence no specialized reconstruction techniques are employed beyond the PF method.

Hadronic tau lepton decays (τ_h), on the other hand, are reconstructed in CMS using the dedicated Hadrons-Plus-Strips (HPS) algorithm [174–176]. The main challenge lies in efficiently differentiating τ_h candidates from the overwhelming quark and gluon jet background in pp collisions. Nonetheless, hadronic tau decays are typically characterized by lower particle multiplicity and more collimated and isolated decay signatures compared to multijet events, thus providing crucial discriminative features that are exploited by the HPS algorithm.

The HPS algorithm uses PF jets, reconstructed with the anti-k_T algorithm with $\Delta R = 0.4$, as seeds. It then assesses whether particles within these jets can be categorized as charged hadrons or π^0 candidates. Identification of π^0 candidates is crucial as the majority of hadronic tau decays contain at least one π^0 in the final state, as depicted in Tab. 4.2. The neutral pions decay almost exclusively into photon pairs. As these photons pass through the tracker material, they often convert into e^+e^- pairs, which get spatially separated in the η-ϕ plane by the intense magnetic field of the CMS solenoid. Due to this separation, the clusters formed in the ECAL by the π^0 decay products typically exhibit an elongated *strip-like* shape in the ϕ direction.

In order to reconstruct the $\pi^0 \to \gamma\gamma$ decays, the HPS algorithm follows a systematic, iterative approach to cluster the photon and electron constituents of the τ_h-seeding jet into *strips* in the η-ϕ plane. It starts by selecting the electron or photon with the highest p_T, not yet assigned to any strip, to seed a new strip. The η and ϕ values of this new e/γ seed are set as the initial position of the strip. Subsequently, the algorithm searches for the e/γ candidate with the next highest p_T within a predefined $\Delta\eta$-$\Delta\phi$ window centered on the strip position. The dimensions of this strip window are dynamically allocated[2] based on the p_T of the strip and the considered e/γ candidate, as described in Ref. [176]. If a suitable candidate is identified within this window, it is merged into the strip, and the position of the strip is recomputed utilizing the p_T-weighted average of the η and ϕ values of all the strip constituents. This procedure repeats until no more electrons or photon candidates are found within the $\Delta\eta$-$\Delta\phi$ window, after which the procedure starts anew for a new strip.

Following the strip reconstruction, hadronic tau candidates are constructed by combining these strips with the charged PF candidates from the τ_h-seeding jet. Using a combinatorial approach, multiple decay mode hypotheses are considered for each jet [177]. These hypotheses correspond to the main tau decay modes listed in Tab. 4.2 and are formed by combining one, two, or three charged particles with up to two strips. Additional constraints are then imposed on the reconstructed τ_h candidates to ensure compatibility with a genuine τ_h. For the decay mode hypotheses corresponding to a resonance-mediated decay, the mass of the τ_h candidate is required to be loosely compatible with that of the resonance. The charge of each τ_h candidate is required to be ± 1. All reconstructed $h^\pm$ and π^0 need to lie within a signal cone, with radius $\Delta R = 3.0/p_T(\tau_h)(\text{GeV})$ (limited to the range 0.05–0.1) with respect to the τ_h momentum. Finally, among the τ_h candidates selected for each seeding jet, the one with the highest p_T is chosen.

For genuine τ_h candidates with $p_T > 20\,\text{GeV}$ and $|\eta| < 2.3$, the HPS algorithm maintains a misreconstruction rate between 11% for $h^\pm$ and 25% for $h^\pm\pi^0$ decays. The overall reconstruction efficiency is mostly limited by a 90% efficiency in charged-hadron track reconstruction [177]. The charge assignment is 99% accurate in the decay mode hypotheses without missing charged hadrons for an inclusive $Z \to \tau\tau$ sample, 98% for τ_h with $p_T \approx 200\,\text{GeV}$, and 92% for τ_h with $p_T \approx 1\,\text{TeV}$.

After reconstructing τ_h candidates, CMS employs a DNN-based identification algorithm called DeepTau to efficiently distinguish genuine tau leptons from misidentified jets, electrons, or muons [177]. The algorithm processes a comprehensive set of input features, including particle candidate properties (p_T, η, ϕ, energy, and particle type), tau-specific characteristics (decay mode, isolation, and signal cone properties), and global event information (pileup density and vertex positions) to differentiate taus from light leptons and jets. Its sophisticated architecture, incorporating convolutional and dense layers, allows for tau identification with higher efficiency and lower misidentification rates compared to the previously used cut-based or simpler machine-learning methods.

[2] This flexibility is essential as charged pions from τ_h decays may undergo nuclear interaction with the tracker material and produce secondary particles with lower p_T, which might escape a fixed-size strip window. Additionally, electron pairs from π^0 decays may end up outside a fixed-size window due to scattering and bremsstrahlung effects. Conversely, high-p_T τ_h decays result in more collimated decay products, allowing the use of smaller strip sizes to reduce background contributions effectively.

The $H \to a_1 a_1$ analysis presented in this thesis does not employ the conventional τ_h reconstruction algorithm, for reasons explained later in Sec. 5.2. Nonetheless, a description of this method is included here for completeness.

4.2.8 Missing Transverse Energy

Though the CMS detector is capable of detecting most stable and long-lived particles, neutrinos and other hypothetical neutral weakly interacting particles pass straight through the detector and evade direct detection. Their presence is inferred indirectly from an imbalance in the recorded energy in the detector's transverse plane. As discussed in Sec. 3.2.1, momentum projected onto the transverse plane is conserved. In proton-proton collisions, the colliding partons carry negligible transverse momentum relative to the beam direction, ensuring that the initial transverse momentum of the system is effectively zero. The missing transverse momentum, $\vec{p}_T^{\,\text{miss}}$, can therefore be defined as the negative vectorial sum of the transverse momenta of all final-state particles in an event:

$$\vec{p}_T^{\,\text{miss}} = -\sum_{i=0}^{N} \vec{p}_T^{\,i}. \tag{4.9}$$

CMS employs two $\vec{p}_T^{\,\text{miss}}$ reconstruction algorithms, both based on PF candidates [178, 179]:

- The PF $\vec{p}_T^{\,\text{miss}}$ is defined exactly as in Eq. 4.9, where the negative vector sum operates on all PF candidates in the event.
- The PUPPI $\vec{p}_T^{\,\text{miss}}$ incorporates pileup mitigation by assigning weights, w_i, to PF candidates based on their likelihood of originating from the primary vertex. Candidates likely associated with pileup are given lower weights. The weighted $\vec{p}_T^{\,\text{miss}}$ is computed as:

$$\vec{p}_T^{\,\text{miss}} = -\sum_{i=0}^{N} w_i \vec{p}_T^{\,i}. \tag{4.10}$$

Both methods are influenced by factors such as the non-linear calorimeter response to hadrons, calorimeter energy thresholds, and track reconstruction inefficiencies. To reduce biases, the so-called "type-1" corrections are applied, which adjust $\vec{p}_T^{\,\text{miss}}$ using jet energy corrections as follows:

$$\vec{p}_T^{\,\text{miss,corr}} = \vec{p}_T^{\,\text{miss}} - \sum_{\text{jets}} (\vec{p}_{T,\text{jet}}^{\,\text{corr}} - \vec{p}_{T,\text{jet}}) \tag{4.11}$$

where the sum runs over all jets with $p_T > 15\,\text{GeV}$ and $\vec{p}_{T,\text{jet}}^{\,\text{corr}}$ and $\vec{p}_{T,\text{jet}}$ denote the energy-corrected and uncorrected jets.

Most CMS analyses use the standard type-1 corrected PF $\vec{p}_T^{\,\text{miss}}$. However, analyses sensitive to pileup effects often employ PUPPI $\vec{p}_T^{\,\text{miss}}$ for improved accuracy, benefiting from its enhanced pileup suppression.

Chapter 5

Search for Light Pseudoscalars in Exotic Decays of the Higgs boson

Contents

5.1 Motivation and Overview 82
5.2 Analysis Strategy 84
5.3 Data and Simulated Samples 85
5.4 Triggers 88
5.5 Offline Event Selection 92
5.5.1 Primary Vertex Selection, Noise Filters, and HEM Veto 92
5.5.2 Veto on b-tagged Jets 92
5.5.3 Muon Selection 94
5.5.4 Track Selection 95
5.5.5 Isolation Requirements 97
5.5.6 Selection of a_1 Candidates 98
5.5.7 Final Selection 98
5.6 Corrections to Simulations 100
5.6.1 Pileup Reweighting 100
5.6.2 Prefiring Weights 100
5.6.3 Tracking Efficiency Corrections 100
5.6.4 Muon Momentum Scale Corrections 100
5.6.5 Muon Identification Efficiency Corrections 100
5.6.6 Trigger Efficiency 101
5.6.7 Track Isolation and Identification Efficiency 104
5.6.8 Higgs-p_T Reweighting 106
5.6.9 Jet Energy Corrections and b-jet Identification Efficiency 106
5.7 Signal Extraction Stategy 106
5.8 Signal Modeling 108
5.9 Background Modeling 112
5.9.1 Modeling of $f_{1D}(i)$ 113

5.9.2 Modeling of $C(i, j)$ 114
5.10 Systematic Uncertainties 122
5.10.1 Uncertainties affecting background 122
5.10.2 Uncertainties affecting signal distributions 123

As previously mentioned in Chapter 2, the success of the standard model (SM) of particle physics comes with inherent limitations, which invite significant exploration into beyond-standard-model (BSM) theories. One intriguing possibility is that the observed 125 GeV Higgs boson, henceforth labeled as H, may be part of an extended scalar sector—a concept supported by several theoretical models. Particularly captivating are models that predict the existence of a light pseudoscalar boson a_1, which could interact with the SM-like Higgs boson. If the a_1 boson is sufficiently light, it could lead to exotic decays of the form $H \to a_1 a_1$. The search for $H \to a_1 a_1$ decays, therefore, provides an effective probe for possible extensions of the SM and the discovery of new physics phenomena. These searches are further warranted by the significant branching ratio still available for Higgs boson decays into BSM particles.

After laying the groundwork with the necessary theoretical and experimental aspects in earlier chapters, this chapter is dedicated to the search for light pseudoscalar bosons produced in exotic decays of the 125 GeV Higgs boson in the 4τ final state. The chapter describes the tailored approach adopted to target the boosted event topologies typical for very light bosons. Critical aspects of the analysis, including the event selection process, signal modeling, and the dedicated data-driven background estimation techniques employed, are discussed in detail. The preliminary results of this analysis have been released as a CMS public note [18]. At the time of writing this thesis, final preparations for journal publication are underway.

5.1 Motivation and Overview

As established by now, a wide variety of beyond-standard-model (BSM) theories, including the Two-Higgs-Doublet Model with an additional scalar singlet (2HDM+S), predict an extended scalar sector that encompasses the well-known 125 GeV Higgs boson and a light pseudoscalar boson a_1. The direct production of these light pseudoscalar bosons at colliders like the LHC is often suppressed due to their weak coupling to standard model (SM) particles. This makes their detection challenging amid the LHC's typical collision products. However, the decay $H \to a_1 a_1$ becomes kinematically feasible if the mass of the pseudoscalar $m_{a_1} < \frac{m_H}{2}$. This decay channel allows the Higgs boson, which is abundantly produced, to serve as a probe for new physics. Existing measurements of Higgs boson couplings permit a substantial branching fraction for decays to BSM particles, with the ATLAS and CMS experiments setting upper limits at 95% confidence level (CL) of 12% and 16%, respectively; this further motivates the searches for these decays [9, 10].

The a_1 boson can decay into a variety of SM particles, as detailed in Chapter 2, thus offering a rich phenomenology for $H \to a_1 a_1$ decays. Consequently, several searches have been carried out by the ATLAS and CMS Collaborations, exploring different decay modes of the a_1 bosons (including 4e, 4μ, $2e\mu$, $2\mu 2\tau$, 4γ, $2b2\mu$, $2b2\tau$, 4b, $\gamma\gamma jj$), covering the mass range $0.2\,\text{GeV} < m_{a_1} < 62.5\,\text{GeV}$ [180–196].

For a_1 bosons with a mass less than 10 GeV (i.e., below the b-quark pair production threshold), leptonic decay channels, particularly $a_1 \to \tau\tau$, are of special relevance as these decays are dominant in this range. Additionally, in specific models such as the Type-III 2HDM+S with high $\tan\beta$, this channel can dominate across the entire (kinematically allowed) mass range of the a_1 boson.

The study presented in this thesis primarily focuses on a search for very light bosons in the $H \to a_1a_1 \to 4\tau$ decay channel, supplemented by potential events from the $H \to a_1a_1 \to 2\mu2\tau$ channel. It covers pseudoscalar masses ranging from 4 to 15 GeV. The analysis is based on proton-proton (pp) collision data at a center-of-mass energy of 13 TeV, corresponding to an integrated luminosity of 138 fb^{-1}, recorded by the CMS detector between 2016 and 2018. This analysis updates a previous CMS search performed with 35.9 fb^{-1} of data collected at 13 TeV. It set 95% CL upper limits between 0.022 and 0.23 on the Higgs boson production cross section times the branching fraction for $H \to a_1a_1 \to 4\tau$ relative to the SM Higgs production cross section, $\sigma\mathcal{B}(H \to a_1a_1 \to 4\tau)/\sigma_{SM}$ [197]. The results of this analysis are displayed in Fig. 5.1.

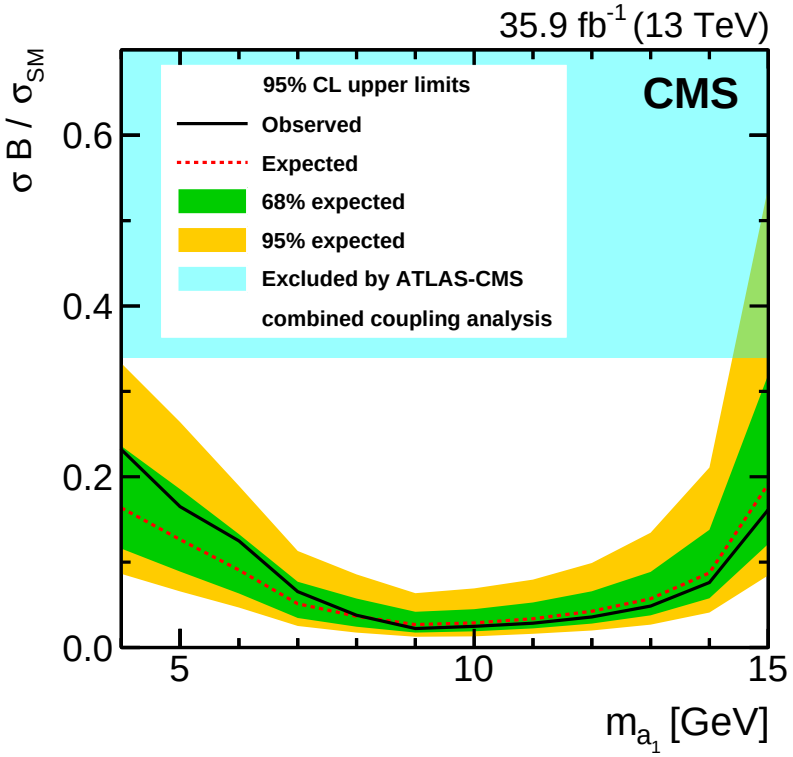

Figure 5.1: The observed and expected upper limits at 95% confidence level on the product of signal cross section and the branching fraction $H \to a_1a_1 \to 4\tau$, relative to the SM inclusive Higgs boson production cross section, set using 35.9 fb^{-1} data collected at 13 TeV by the CMS experiment [197].

The current study significantly extends the previous analysis by utilizing a data set approximately four times larger, benefitting from state-of-the-art calibrations provided for Run 2. Additionally, tighter selection criteria and a b-jet veto have been imposed to effectively suppress background events.

Initially, the objective was to analyze only the 2018 data and combine it with the previously published results from 2016. However, the performance improvements achieved with the new analysis strategy exceeded those expected from the increase in luminosity alone. This led to the decision to reanalyze the 2016 data using this more effective strategy. Also, the 2017 data was initially excluded due to the absence of the nominal trigger required for the analysis. However, extensive studies identified an alternative trigger with sufficient sensitivity, enabling the inclusion of the 2017 data in the study. As a result of these concerted efforts throughout the duration of the author's PhD, the study now encompasses the entire luminosity of Run 2, thus achieving increased sensitivity for the search.

5.2 Analysis Strategy

As mentioned in the previous section, this search focuses on light pseudoscalars with m_{a_1} ranging from 4 to 15 GeV. Due to the significant mass difference between the a_1 bosons and the Higgs boson, the a_1 bosons produced in the $\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1$ decays are given a large Lorentz boost. As a result, the decay products of a_1 are highly collimated, giving rise to overlapping-fermion signatures in the CMS detector.

The highly collimated nature of these decays significantly complicates the identification of $\mathrm{a}_1 \to \tau\tau$ decays. The Hadron Plus Strips (HPS) algorithm, discussed in Section 4.2.7, struggles to efficiently identify individual hadronic tau decays under these conditions. The close proximity of energy deposits in the calorimeter and adjacent tracks in the tracker often leads to the misidentification of τ decay modes, ineffective isolation of signals, and potential misassignment of energies to the decay products. Furthermore, the standard tau triggers in CMS, which require high p_{T} thresholds to reject the abundance of events produced by quantum chromodynamics (QCD) multijet processes, are not ideally suited for detecting these closely spaced, low-energy tau decay events.

To address these challenges and effectively identify highly Lorentz-boosted tau lepton pairs with overlapping decays, this analysis employs a specialized strategy, deviating from standard reconstruction and selection techniques. In each a_1 decay leg, one of the tau leptons is identified through its muonic decay, leveraging the CMS detector's excellent muon reconstruction capabilities. The other tau lepton is allowed to decay into a lepton or a single charged hadron with additional neutral particles, as illustrated in Figure 5.2. This "1-prong" decay of the tau is identified by the presence of a single, oppositely-charged track close to the muon[1].

This strategy also supports the identification of $\mathrm{a}_1\mathrm{a}_1 \to (\mu\mu)(\tau_\mu\tau_{\text{1-prong}})$ decays, which are, hence, included for further analysis. However, the $\mathrm{a}_1\mathrm{a}_1 \to 4\mu$ decay, which could potentially present a similar topology, is not considered due to its very low event yields and minimal contribution to the study's sensitivity.

This analysis predominantly targets the dominant gluon-gluon fusion (ggF) Higgs production mechanism, typically producing the H boson with relatively low transverse momentum (p_{T}). As a result, the two a_1 bosons are emitted nearly back-to-back in the transverse plane, leading to a substantial separation in azimuthal angle ($\Delta\phi$) between their decay products. When the

[1] Note that 3-prong tau decays are excluded due to their high QCD multijet background and lower reconstruction efficiency.

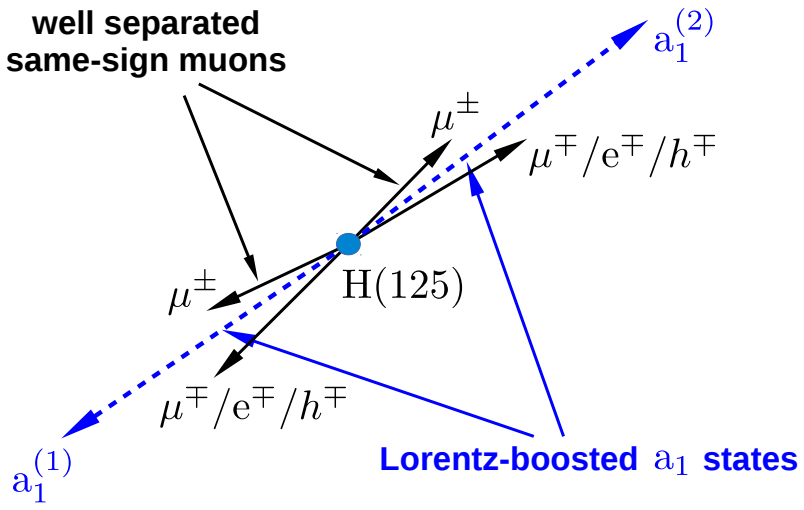

Figure 5.2: Illustration of the signal topology.

H boson is produced with high p_{T}, such as from initial state radiation in ggF processes or in other production mechanisms like vector boson fusion (VBF) or associated productions (VH or $\mathrm{t\bar{t}H}$), the $\Delta\phi$ between the $\mathrm{a_1}$ bosons decreases, though their separation in pseudorapidity ($\Delta\eta$) may still be significant.

So, the analysis strategically focuses on identifying same-sign (SS) dimuon events with significant angular separation, $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}$, where each muon is accompanied by a nearby oppositely charged particle from the same $\mathrm{a_1}$ decay. The same-sign criterion reduces background from processes such as top quark pair production ($\mathrm{t\bar{t}}$), Drell–Yan, and diboson production, thereby enhancing the sensitivity of the search for these rare decay events. The two-dimensional distribution of the invariant masses of the two muon-track systems, $\mathrm{2D}(m_1, m_2)$, is used to discriminate the signal from the background. The distribution shapes for the signal processes are obtained using dedicated simulated samples. A data-driven approach, employing control regions, is utilized to estimate the shape of the background processes. All of these steps are thoroughly detailed in the subsequent sections. Ultimately, a simultaneous signal-plus-background likelihood fit is applied to the $\mathrm{2D}(m_1, m_2)$ distribution to extract $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1} \to 4\tau)/\sigma_{\mathrm{SM}}$. Further details on the signal extraction procedure will be discussed in the next chapter.

5.3 Data and Simulated Samples

The data utilized in this analysis was recorded by the CMS detector during Run 2 (2016–2018) at a center-of-mass energy of 13 TeV using double muon high-level triggers. The collected data are classified into the DoubleMuon primary dataset, which is further divided into subsets referred to as *eras*. These eras are defined based on notable changes in LHC operations and detector conditions. Detailed information about the DoubleMuon dataset, including the

corresponding eras, run ranges, and luminosities, is outlined in Table 5.1. The total integrated luminosity collected during Run 2 amounts to $138\,\mathrm{fb}^{-1}$, spread across the years as follows: $36.3\,\mathrm{fb}^{-1}$ in 2016, $41.5\,\mathrm{fb}^{-1}$ in 2017, and $58.9\,\mathrm{fb}^{-1}$ in 2018. For the purposes of this analysis, only data collected under stable beam conditions with all subdetectors fully operational and certified by CMS as suitable for physics analyses were included.

Dataset	Run range	Luminosity [fb^{-1}]
	2016	
`2016B-ver2`	273150–275376	5.83
`2016C`	275656–276283	2.60
`2016D`	276315–276811	4.29
`2016E`	276831–277420	4.07
`2016F`	277932–278808	3.44
`2016G`	278820–280385	7.65
`2016H`	281613–284044	8.74
	2017	
`2017B`	297047–299329	4.80
`2017C`	299368–302029	9.57
`2017D`	302031–302663	4.25
`2017E`	303824–304797	9.32
`2017F`	305040–306462	13.54
	2018	
`2018A`	315257–316995	14.03
`2018B`	317080–319310	7.07
`2018C`	319337–320065	6.89
`2018D`	320500–325175	31.84

Table 5.1: Overview of the datasets used in the analysis, listing run ranges and integrated luminosities by era from 2016 to 2018.

The Monte Carlo (MC) simulated events for the $\mathrm{H} \to \mathrm{a_1 a_1} \to 4\tau$ process are produced at leading order (LO) using the PYTHIA (v.8.212) event generator [131, 132] for four major Higgs boson production mechanisms, namely: gluon-gluon fusion (ggF), vector boson fusion (VBF), vector boson associated production (VH), and top quark pair associated production ($\mathrm{t\bar{t}H}$). Additionally, the samples of the $\mathrm{H}(125) \to \mathrm{a_1 a_1} \to 2\mu 2\tau$ process are generated at LO for the ggF Higgs production mode using MADGRAPH5_aMC@NLO (v.2.6.5) [133]. The MADGRAPH5_aMC@NLO generator is interfaced with PYTHIA to simulate parton showering and hadronization. The simulations for both processes cover $\mathrm{a_1}$ masses from 4 to 15 GeV in 1 GeV increments. Details of the simulated signal samples are presented in Table 5.2.

Table. 5.3 lists the different simulated background processes, along with their cross sections at 13 TeV. QCD multijet production, which forms the primary background source, is modeled using PYTHIA (v.8.212) in $\hat{p_T}$-binned samples. The produced events are muon-enriched, meaning they contain at least one muon with transverse momentum exceeding 5 GeV. The W/Z+jets process with subsequent leptonic decays of W and Z bosons is simulated at leading order with MADGRAPH5_aMC@NLO (v.2.6.5) and interfaced with PYTHIA for parton showering and hadronization. The inclusive single top production and top-pair production decaying via the leptonic, semi-leptonic, and hadronic channels are generated using the POWHEGBOX

Dataset	a_1 mass range [GeV]	MC generator
ggF, $H \to a_1a_1 \to 4\tau$	4–15	PYTHIA
VBF, $H \to a_1a_1 \to 4\tau$	4–15	PYTHIA
VH, $H \to a_1a_1 \to 4\tau$	4–15	PYTHIA
$t\bar{t}H$, $H \to a_1a_1 \to 4\tau$	4–15	PYTHIA
ggF, $H \to a_1a_1 \to 2\mu2\tau$	4–15	MADGRAPH5_aMC@NLO

Table 5.2: List of the simulated datasets for the $H \to a_1a_1 \to 4\tau$ and $H(125) \to a_1a_1 \to 2\mu2\tau$ processes, along with the MC generators used to produce them. Each dataset is generated for a_1 mass ranging from 4 to 15 GeV in 1 GeV increments.

Dataset	MC generator	Cross section [pb]
QCD μ-enriched, $\hat{p_T}$ = 20-30 GeV	PYTHIA	558528000*0.0053 (LO)
QCD μ-enriched, $\hat{p_T}$ = 30-50 GeV	PYTHIA	139803000*0.01182 (LO)
QCD μ-enriched, $\hat{p_T}$ = 50-80 GeV	PYTHIA	19222500*0.02276 (LO)
QCD μ-enriched, $\hat{p_T}$ = 80-120 GeV	PYTHIA	2758420*0.03844 (LO)
QCD μ-enriched, $\hat{p_T}$ = 120-170 GeV	PYTHIA	469797*0.05362 (LO)
QCD μ-enriched, $\hat{p_T}$ = 170-300 GeV	PYTHIA	117989*0.07335 (LO)
QCD μ-enriched, $\hat{p_T}$ = 300-470 GeV	PYTHIA	7820.25*0.10196 (LO)
QCD μ-enriched, $\hat{p_T}$ = 470-600 GeV	PYTHIA	645.528*0.12242 (LO)
QCD μ-enriched, $\hat{p_T}$ = 600-800 GeV	PYTHIA	187.109*0.13412 (LO)
QCD μ-enriched, $\hat{p_T}$ = 800-1000 GeV	PYTHIA	32.3486*0.14552 (LO)
W($\to l\nu_l$)+jets	MADGRAPH5_aMC@NLO	61530 (NLO)
Z($\to ll$)+jets, $m_{ll} > 50$ GeV	MADGRAPH5_aMC@NLO	6077 (NLO)
Z($\to ll$)+jets, $m_{ll} < 50$ GeV	MADGRAPH5_aMC@NLO	18610 (LO)
$t\bar{t} \to 2l2\nu_l$	POWHEGBOX	88.29 (NLO)
$t\bar{t} \to l\nu + q$	POWHEGBOX	365.35 (NLO)
$t\bar{t} \to 2q$	POWHEGBOX	377.96 (NLO)
t + W production	POWHEGBOX	35.6 (NLO)
$\bar{t}$ + W production	POWHEGBOX	35.6 (NLO)
Single top t-channel production	POWHEGBOX	136.95 (NLO)
Single anti-top t-channel production	POWHEGBOX	80.95 (NLO)
WW inclusive	PYTHIA	115.0 (LO)
WZ inclusive	PYTHIA	47.13 (LO)
ZZ inclusive	PYTHIA	16.52 (LO)

Table 5.3: Simulated datasets for background processes used in the analysis, along with the corresponding Monte Carlo generator and cross sections at 13 TeV. For the QCD multijet samples, the product of the cross sections and filter efficiencies is listed. The cross sections are indicated with labels LO (leading order), NLO (next-to-leading order), or NNLO (next-to-next-to-leading order) to indicate the precision of the calculations.

(v.2.0) generator [136–138], which is also interfaced with PYTHIA. Inclusive diboson processes (WW, WZ, ZZ) are generated directly with PYTHIA. Although these background samples are not used directly to estimate background contributions, they are essential for designing and optimizing the signal selection criteria and deriving correction factors, as discussed later in this chapter.

Both the simulated signal and background samples use the CP5 tune [130] for modeling underlying events and employ the NNPDF3.1 parton distribution functions [125]. For samples generated with MADGRAPH5_aMC@NLO, the MLM matching scheme is used for LO samples [134], while the FxFx scheme is applied for next-to-leading order (NLO) samples [135].

Distinct MC samples are produced for each data-taking year in order to account for the specific operating conditions, detector configurations, and calibrations during those periods.

5.4 Triggers

As outlined in Sec. 3.2.2.6, triggers play a crucial role in determining which events are recorded for further analysis among the vast quantity of data generated by collisions at the LHC. This analysis employs a set of double muon triggers, listed in Table 5.4, in order to select events of interest. In 2016 and 2018, same-sign dimuon triggers were used, which, as the name suggests, select pairs of muons with the same electric charge. These triggers do not impose any isolation requirements on muons. In 2016, the trigger mandates that the p_{T} of the leading and subleading muon exceed 17 GeV and 8 GeV, respectively. In 2018, the requirements were raised to 18 GeV for the leading muon and 9 GeV for the subleading muon. Additionally, one of the triggers in 2016 also requires that the tracks of the two muons have their points of closest approach to the beam axis within 2 mm of each other along the longitudinal direction.

Year	HLT trigger path	Luminosity [fb^{-1}]
2016	`HLT_Mu17_Mu8_SameSign_DZ`	11.0
	`HLT_Mu17_Mu8_SameSign`	25.3
2017	`HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass8`	41.5
2018	`HLT_Mu18_Mu9_SameSign`	59.8

Table 5.4: HLT triggers and the corresponding integrated luminosities used in the analysis.

In 2017, due to technical limitations, the trigger imposing the same-sign criterion was not available for the full data-taking period. Hence, a dimuon trigger that imposes loose isolation on muons is used instead as an alternative. This loose isolation requirement is defined as:

$$I_{\mu} = \frac{\sum p_{\mathrm{T}}(\text{charged hadrons})}{p_{\mathrm{T}}(\mu)} < 0.4, \tag{5.1}$$

where the p_{T} sum in the numerator runs over all charged hadrons within an isolation cone of 0.3 around the muon. In addition, the trigger applies p_{T} thresholds of 17 GeV and 8 GeV on the leading and subleading muon.

Despite initial expectations that the isolation requirement might reduce the selection efficiency of signal events, wherein muons are accompanied by close-by tracks from the same

$a_1 \rightarrow \tau_\mu \tau_{1-\text{prong}}$ decay leg, the trigger demonstrated relatively high efficiency in selecting these events. This can be observed, for example, in the left plot in Fig. 5.3 that compares the selection efficiencies of the double muon trigger with loose isolation requirement to that of the same-sign double muon trigger without any isolation requirement. This comparison is based on the simulated $gg \rightarrow H \rightarrow a_1 a_1$ signal sample for 2017. It is important to note, for clarity, that even though the same-sign trigger was deployed late during data taking in 2017, it was included in the simulation, enabling this comparison. In both scenarios, the full offline selection criteria (detailed in the subsequent section) are applied. The comparison shows only a marginal deterioration in signal efficiency due to the online isolation requirements.

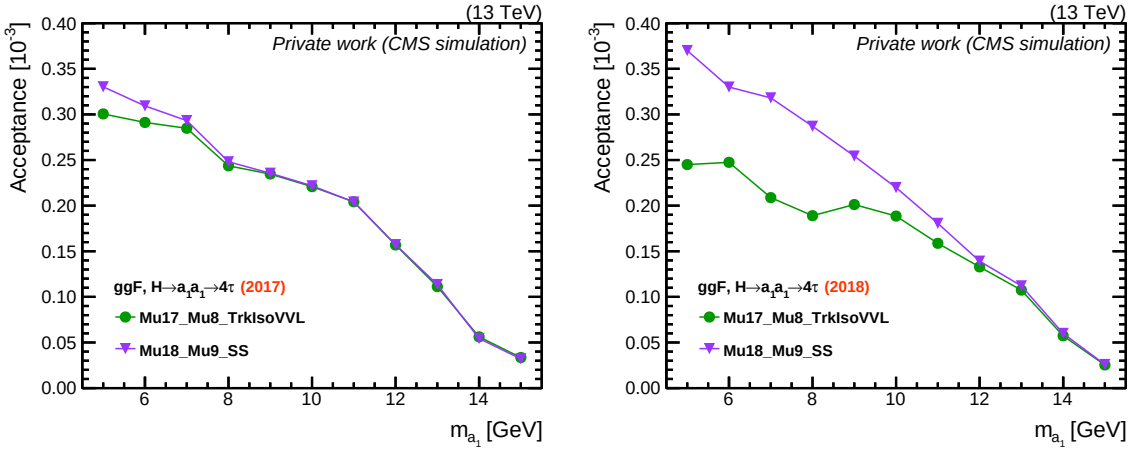

Figure 5.3: Comparison of the acceptance of the $gg \rightarrow H \rightarrow a_1 a_1$ signal between the double muon trigger with loose isolation requirement (green) and double muon trigger with same-sign requirement (violet). The acceptance is computed as a function of m_{a_1}. The left plot illustrates the trigger performances in the 2017 MC samples, whereas the right plot examines the 2018 samples. Only a marginal degradation of acceptance due to isolation requirement is observed in the 2017 samples, whereas a sizable reduction of efficiency is found for 2018.

As a verification, the same check was conducted with the simulated samples for 2018 to understand the performance of the trigger. The right plot in Figure 5.3 shows a more significant efficiency loss for the alternative double muon trigger (as expected), indicating differences in trigger performance across the years.

To further validate these findings, the distribution of the isolation variable I_μ defined in Eq. 5.1 was computed offline for the $gg \rightarrow H \rightarrow a_1 a_1 \rightarrow 4\tau$ signal with $m_{a_1} = 5$ GeV. Figure 5.4 displays these distributions for events selected offline both before and after applying the trigger requirements. A clear deterioration of signal efficiency for $I_\mu > 0.4$ is observed in the 2018 sample, whereas the 2017 samples show no significant impact.

As a final cross-check, the yields in the data sample of same-sign muons selected with the trigger under investigation are compared between the 2017 and 2018 datasets. Figure 5.5 presents the p_T distributions for the leading and trailing muons across these datasets, with no isolation requirement applied offline. Remarkably, the yield of same-sign muon pairs in the 2017 dataset is approximately 1.5 times higher than that in the 2018 dataset, despite the integrated luminosity in 2017 being roughly 1.4 times lower than in 2018. This discrepancy

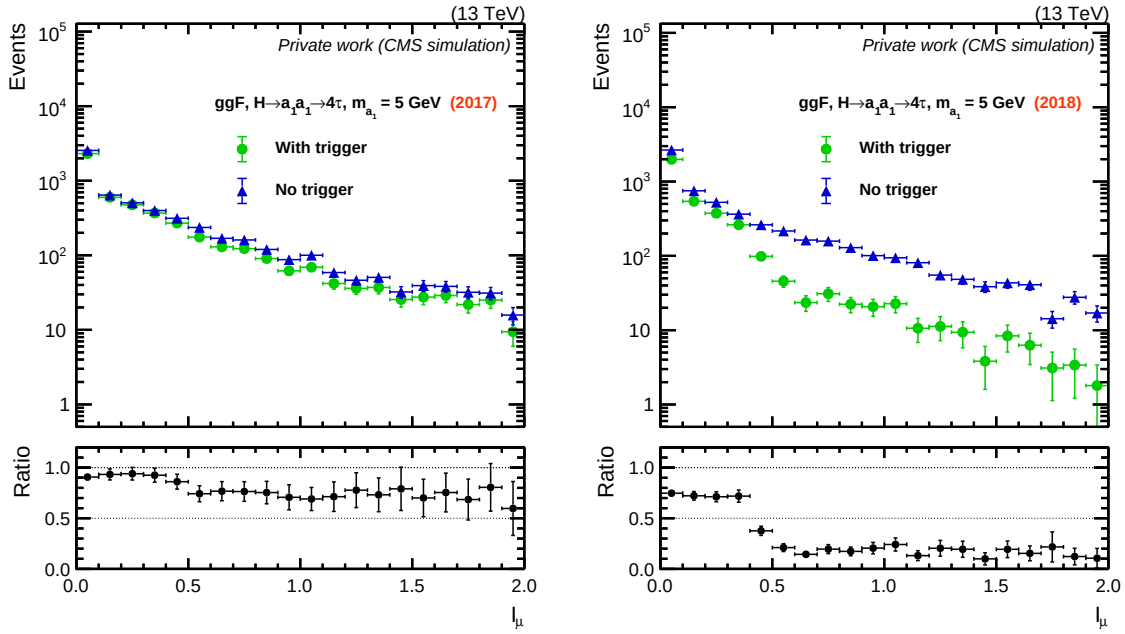

Figure 5.4: Isolation variable I_μ computed offline in the sample of $gg \to H \to a_1 a_1 \to 4\tau$ signal events with $m_{a_1} = 5\,\text{GeV}$ before (blue triangles) and after (red circles) requiring events to pass double muon trigger with loose isolation requirement. The left plot shows the distributions obtained for the 2017 MC sample, and the right plot shows the distributions obtained for the 2018 sample. No drastic drop in efficiency is observed for $I_\mu > 0.4$ in the 2017 samples, whereas significant degradation occurs for the 2018 sample.

is more pronounced at lower muon p_T values, corresponding to higher values of the relative isolation variable. Given that the offline selections are identical, this suggests a higher trigger efficiency in 2017 compared to 2018. However, when an isolation requirement is incorporated into the offline selection, the yield ratio of selected muon pairs aligns more closely with the ratio of integrated luminosities for 2017 and 2018, as illustrated in Figure 5.6.

Following these comprehensive studies, it was determined that the trigger `HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass8` exhibited different responses to non-isolated muons between the 2017 and 2018 datasets. Unlike the 2018 version, the 2017 trigger demonstrated unexpectedly high efficiency in selecting non-isolated muons, which was advantageous for this analysis. Given these findings, this trigger was deemed suitable for analyzing the 2017 data.

Despite a detailed analysis, the exact cause of the differing responses of the trigger between the 2017 and 2018 datasets could not be conclusively identified. The trigger configurations were largely consistent across both years. However, the possibility of software updates or other undocumented modifications cannot be ruled out.

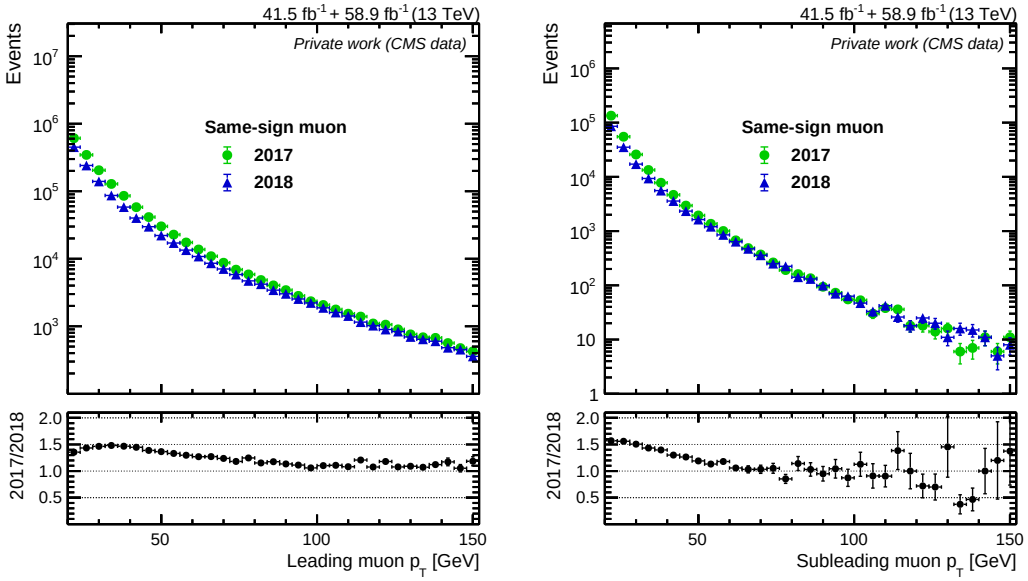

Figure 5.5: Distributions of the leading (left) and subleading (right) muon p_{T} distributions in the data sample of selected same-sign muon pairs. Events are triggered by the double muon trigger with the isolation filter, but no isolation requirement is imposed on muons offline. The yield in the 2017 sample (green) exceeds the yield in the 2018 sample (blue).

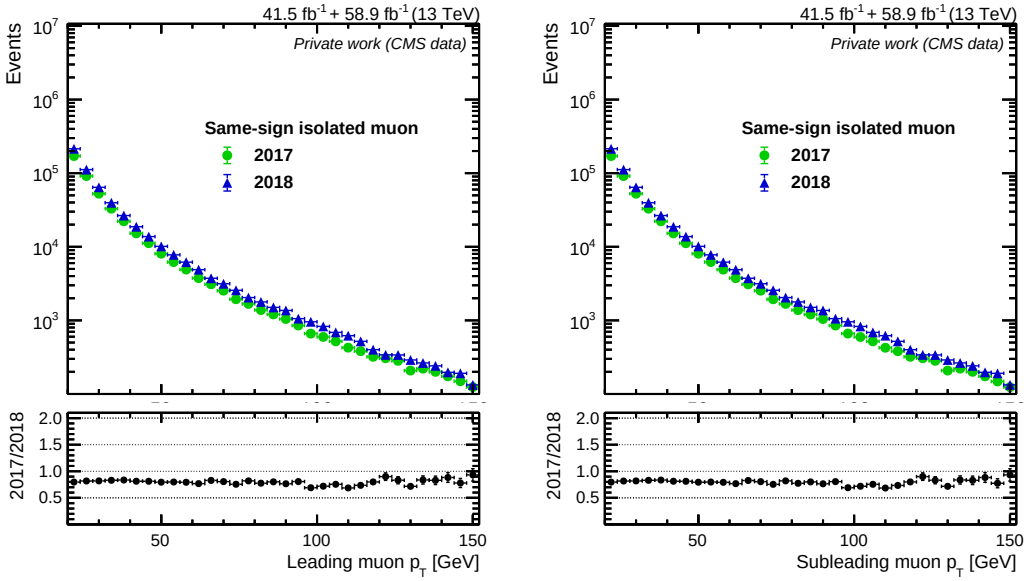

Figure 5.6: Distributions of the leading (left) and subleading (right) muon p_{T} distributions in the data sample of selected same-sign muon pairs. Events are triggered by the double muon trigger with the isolation filter. Isolation requirement of $I_{\mu} < 0.4$ is imposed on muons in offline selection. The ratio of data yields between the 2017 (green) and 2018 (blue) samples approaches the ratio of luminosities collected in the two years.

5.5 Offline Event Selection

Events that meet the criteria set by the triggers are considered for further analysis, which includes offline selection steps designed to effectively reduce the background while maximizing signal retention. This section discusses the specific offline selection steps employed in this analysis. The ensemble of requirements outlined below collectively defines the **signal region (SR)**.

5.5.1 Primary Vertex Selection, Noise Filters, and HEM Veto

Primary Vertex Selection: All selected events are required to contain a well-reconstructed primary vertex corresponding to the hard-scattering process. The identification of the vertex corresponding to the hard scatter is explained in Sec. 4.2.3.

Noise filters: Events are further refined through the application of noise filters, which are essential for eliminating anomalies caused by electronic and detector malfunctions [198].

HEM Veto for 2018: Starting from Run 319077 in 2018, the negative endcap HCAL sectors HEM15 and HEM16 became non-operational. This affected the reconstruction of physics objects in the affected runs in the detector region $-3.0 < \eta < -1.3$ and $-1.57 < \phi < -0.87$. To effectively address this issue, a specific veto policy has been implemented for events recorded during and after Run 319077 [199]. Events are excluded if they contain:

- any electron with $p_T > 30\,\text{GeV}$, $-3.0 < \eta < -1.4$ and $-1.57 < \phi < -0.87$
  or
- any jet with $p_T > 30$ GeV, $-3.2 < \eta < -1.2$, $-1.77 < \phi < -0.67$ and $\Delta\phi(\text{jet}, E_T^{miss}) < 0.5$.

In simulations, events that feature electrons or jets within these kinematic regions are adjusted with a weighting factor corresponding to the ratio of the integrated luminosity of unaffected runs to the total integrated luminosity of 2018. It has been verified that these adjustments have minimal impact on the overall analysis, with a reduction of 0.5% observed in the total yield in data and between 0.2% and 1.0% in the signal yield for 2018.

5.5.2 Veto on b-tagged Jets

A significant background in this analysis arises from QCD multijet events that contain heavy-flavored hadrons decaying into muons. To effectively suppress this background, an early step in the event selection involves applying a veto on events containing b-tagged jets. The identification of b jets is performed using the DeepJet algorithm, described in Sec. 4.2.6. Jets are required to meet the criteria established by the *tight* working point, characterized by a misidentification rate of 0.1%. Consequently, any event that contains at least one b-tagged jet with $p_\text{T} > 20\,\text{GeV}$ and $|\eta| < 2.4$ is excluded from further analysis.

The efficacy of the b-jet veto in reducing background is demonstrated in Fig. 5.7, which displays various variables associated with the leading muon. Implementing the veto results in a substantial reduction of the background by over 60%. For comparison, the lower panel of the figure also presents signal distributions, where only minor variations of less than 15% are observed. This slight decrease in signal can be attributed to the misidentification of the a_1

decay products as b jets. Therefore, the b-jet veto is highly effective in eliminating significant background noise while minimally affecting signal acceptance.

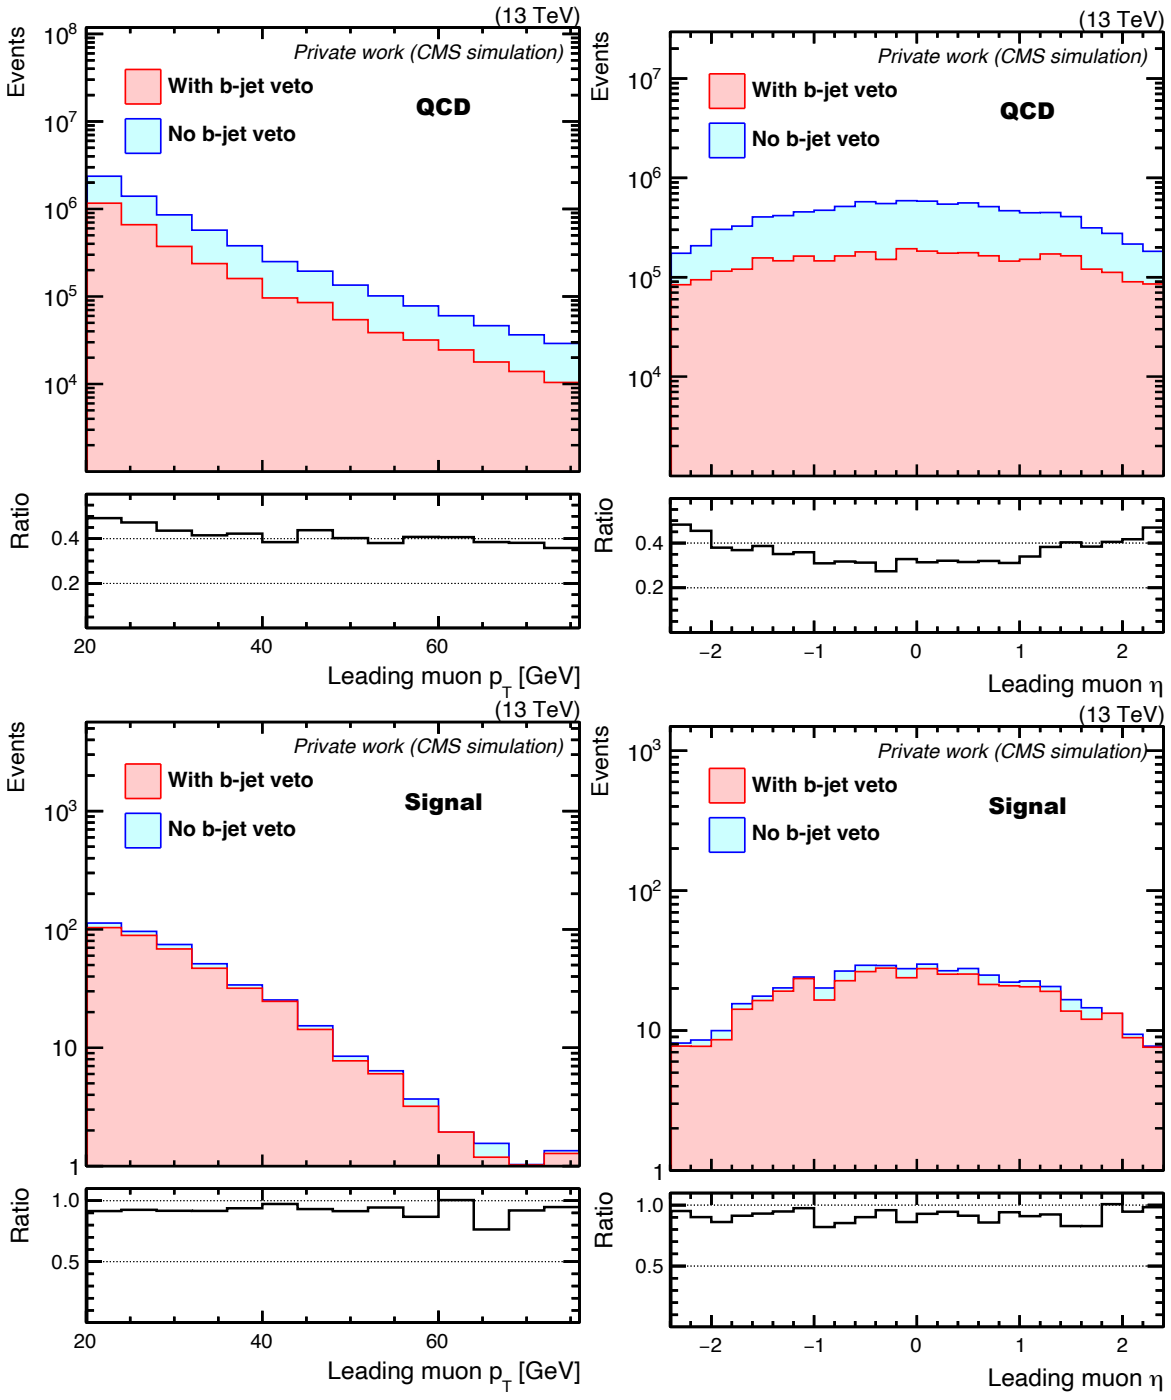

Figure 5.7: Distributions of the p_{T} and η of the leading muon, with and without a b-jet veto. The top panel shows the distributions for the QCD multijet simulated samples. The bottom panel presents the distributions for the ggF $\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1 \to 4\tau$ process, for $m_{\mathrm{a}_1} = 10\,\mathrm{GeV}$.

5.5.3 Muon Selection

The offline selection process proceeds with identifying events containing two particle-flow (PF) muons that pass the *medium* working point of the standard CMS muon identification criteria. The muons are required to satisfy specific kinematic constraints:

- The p_{T} of the leading and subleading muons must be greater than 19 and 10 GeV, respectively.
- Both muons are required to lie within the pseudorapidity region $|\eta| < 2.4$.
- The leading and subleading muons must match with the trigger objects that fired the higher- and lower-p_{T} leg of the trigger, respectively, within a ΔR cone of 0.5.
- The impact parameters of the muons in the transverse and longitudinal plane with respect to the primary vertex must be $|d_0| < 0.05$ cm and $|d_z| < 0.1$ cm, respectively. As detailed in Sec. 4.2.4, the medium muon ID effectively selects not only prompt muons but also those produced in heavy quark decays, which are commonly associated with displaced secondary vertices. Given the significant background contribution of these decays in this analysis, these stringent impact parameter cuts are applied to suppress them.

Upon successfully identifying two muons that meet the above criteria, the pair is required to satisfy the following additional conditions:

- The angular separation between the muons must be $\Delta R > 1.5$. This requirement underscores the substantial separation typically seen between the decay products of the two a_1 bosons. Figure 5.8 illustrates this concept by showing the ΔR distribution between the leading and subleading muons across selected mass points in simulated signal samples.
- Both muons must have the same electric charge, which significantly reduces background contributions from processes like Drell–Yan, top-quark pair production, and diboson processes[2].

If more than one same-sign muon pair is found in an event, the pair with the highest sum of transverse momentum is selected.

Selected kinematic distributions of these muon pairs are showcased in Fig. 5.9, illustrating the observed data and various background processes. It also features the signal distributions for an a_1 mass hypothesis of 10 GeV. The signal is normalized assuming the SM Higgs production cross section and a branching fraction $\mathcal{B}(\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1 \to 4\tau) = 0.05$. This value is chosen conservatively, within the current experimental constraints from CMS that allow up to a 16% branching fraction for Higgs decays to BSM processes [9]. This branching fraction of 0.05 will be consistently used throughout this chapter for all signal illustrations. The distributions in Fig. 5.9 demonstrate that following the application of the above-mentioned selection criteria, QCD multijet events predominantly constitute the background, accounting for approximately 95% of it. Minor contributions to the background also come from processes such as top-antitop ($\mathrm{t\bar{t}}$), single-top, Z+jets, W+jets, and diboson production.

[2]It is important to note that the same-sign requirement is consistently applied during offline selection across all analyzed years, including 2017, regardless of the online trigger criteria.

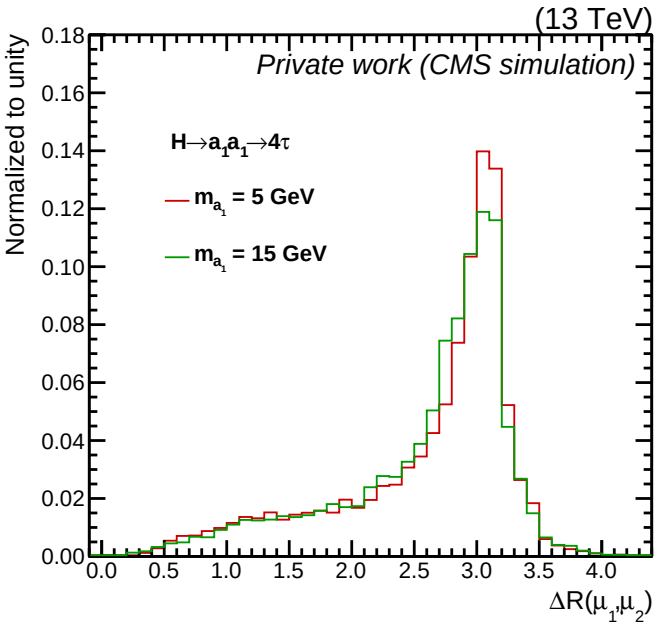

Figure 5.8: Distribution of the angular separation ΔR between the leading and subleading muons in the ggF $H \to a_1 a_1 \to 4\tau$ simulated samples for different a_1 masses.

5.5.4 Track Selection

The analysis also selects *high-purity* tracks associated with reconstructed charged PF objects, excluding the pair of SS muons. These tracks are employed in identifying and isolating the candidates for the $a_1 \to \tau_\mu \tau_{\text{1-prong}}$ or $a_1 \to \mu\mu$ decays, which are hereafter referred to as a_1 candidates. The selection process considers two distinct types of tracks: isolation and signal tracks. The detailed criteria for selecting these tracks and their specific roles in the analysis are outlined in Table 5.5. The signal tracks represent a subset of the isolation tracks, subject to stricter selection criteria.

Track type	p_T	$\lvert\eta\rvert$	$\lvert d_0\rvert$ w.r.t. PV	$\lvert d_z\rvert$ w.r.t. PV	Purpose
Isolation	$> 1.0\,\text{GeV}$	< 2.4	$< 0.2\,\text{cm}$	$< 0.3\,\text{cm}$	Define isolation criterion for the a_1 candidates
Signal	$> 2.5\,\text{GeV}$	< 2.4	$< 0.02\,\text{cm}$	$< 0.04\,\text{cm}$	Build a_1 candidates

Table 5.5: Overview of the types of tracks considered in the analysis, with their selection criteria and objectives.

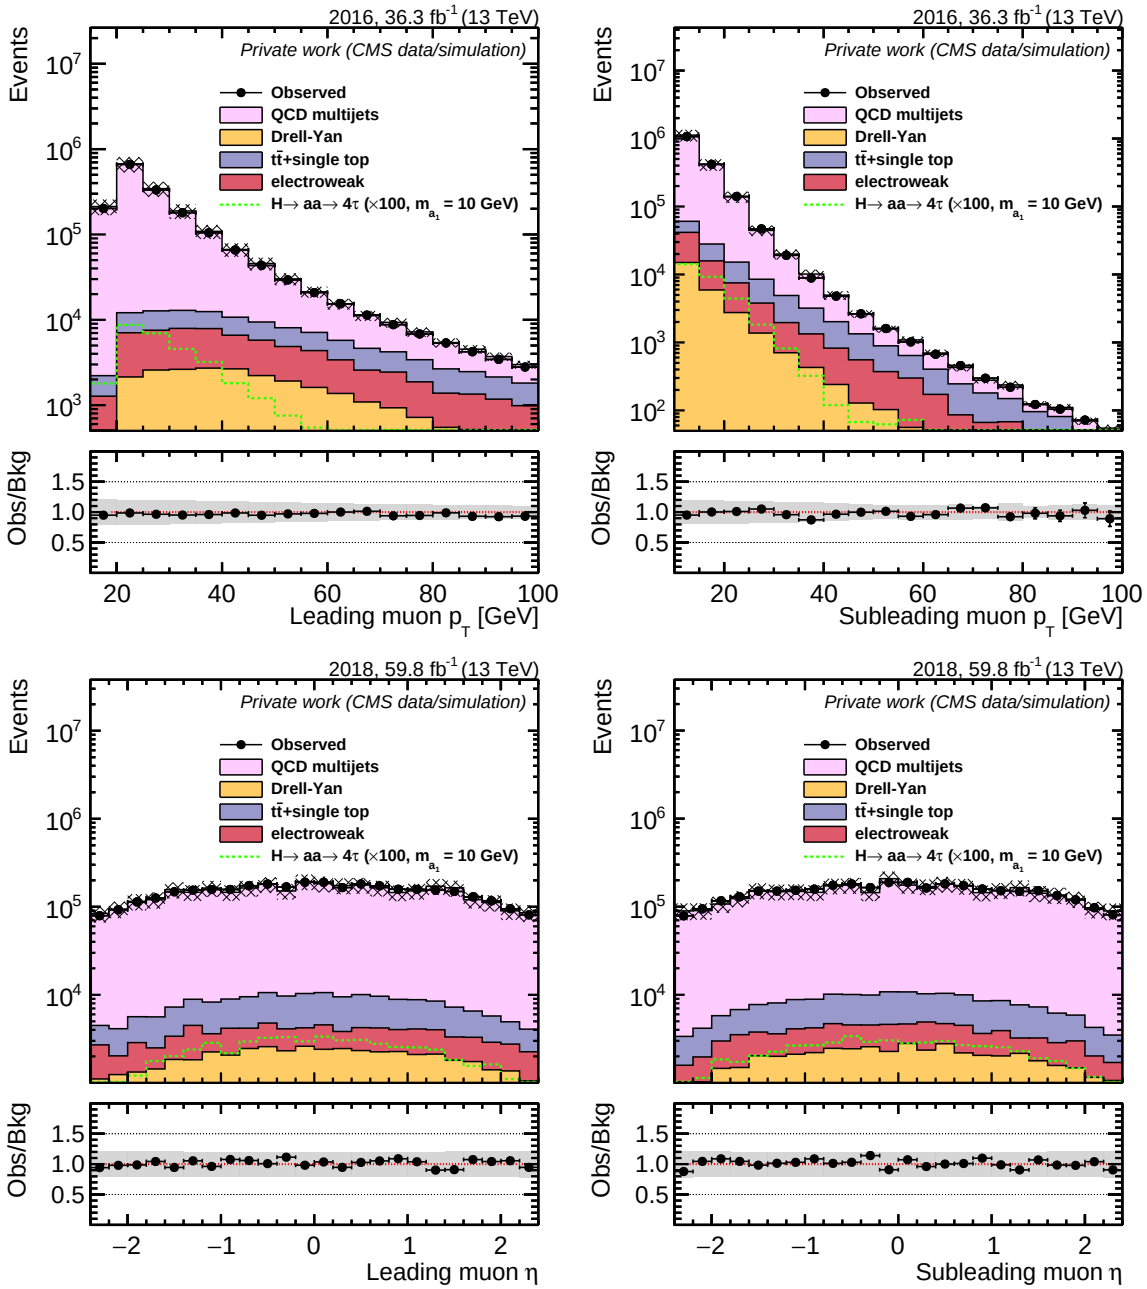

Figure 5.9: Distributions of the leading muon p_{T} (upper left), subleading muon p_{T} (upper right), leading muon pseudorapidity (lower left), and subleading muon pseudorapidity (lower right), in the selected sample of same-sign muon pairs. The data (represented by black dots) are compared with SM background expectations obtained from simulation (solid histograms). Background contributions from W+jets and diboson production processes are collectively aggregated under the "electroweak" category. The dashed histogram illustrates the distribution for the signal in the $\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1 \to 4\tau$ channel for $m_{\mathrm{a}_1} = 10$ GeV. The signal normalization is computed assuming the SM Higgs production cross section and a branching ratio $\mathcal{B}(\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1 \to 4\tau) = 0.05$. The signal is scaled by a factor of 100 for better visualization.

5.5.5 Isolation Requirements

Given the search topology, each muon is expected to have only one nearby track. Accordingly, isolation criteria are implemented wherein each muon from the SS pair is required to have exactly one oppositely-charged isolation track within a ΔR cone of 0.5 around the muon.

The justification for setting the maximum radius of the isolation cone at $\Delta R = 0.5$ comes from a generator-level study of simulated signal samples. The study uses $\mu^{\pm}\pi^{\mp}$ pairs to mimic the reconstructed $\tau_{\mu}\tau_{\text{1-prong}}$ candidates. Figure 5.10 illustrates the angular separation between the muons and charged pions, $\Delta R(\mu^{\pm}, \pi^{\mp})$, for different values of m_{a_1} obtained from this study. As m_{a_1} increases, the decay products of the a_1 boson become less boosted, resulting in greater separation between them. This leads to distinct angular distributions for different m_{a_1} values. For masses ranging from 4 to 15 GeV, $\Delta R = 0.5$ was identified as the optimal cut to ensure exclusive pairing of each muon with its corresponding one-prong track. This cut is particularly ideal for intermediate masses, while it also ensures a good signal-to-background ratio for lower masses and maintains reasonable signal acceptance for higher masses.

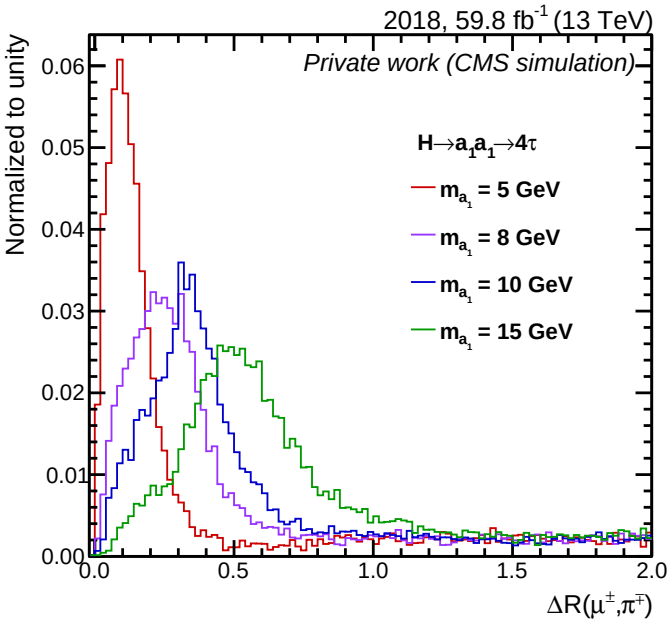

Figure 5.10: Distributions of ΔR of generator-level muon-pion pairs in the $\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1 \to 4\tau$ process, for different a_1 masses.

The stringent isolation requirement described in this section effectively reduces background events. This can be clearly interpreted from Fig. 5.11, which displays the multiplicity of isolation tracks within a $\Delta R = 0.5$ cone around each of the two muons. Background processes, such as the QCD multijet events, exhibit a higher track multiplicity compared to the signal distribution, which characteristically peaks at one.

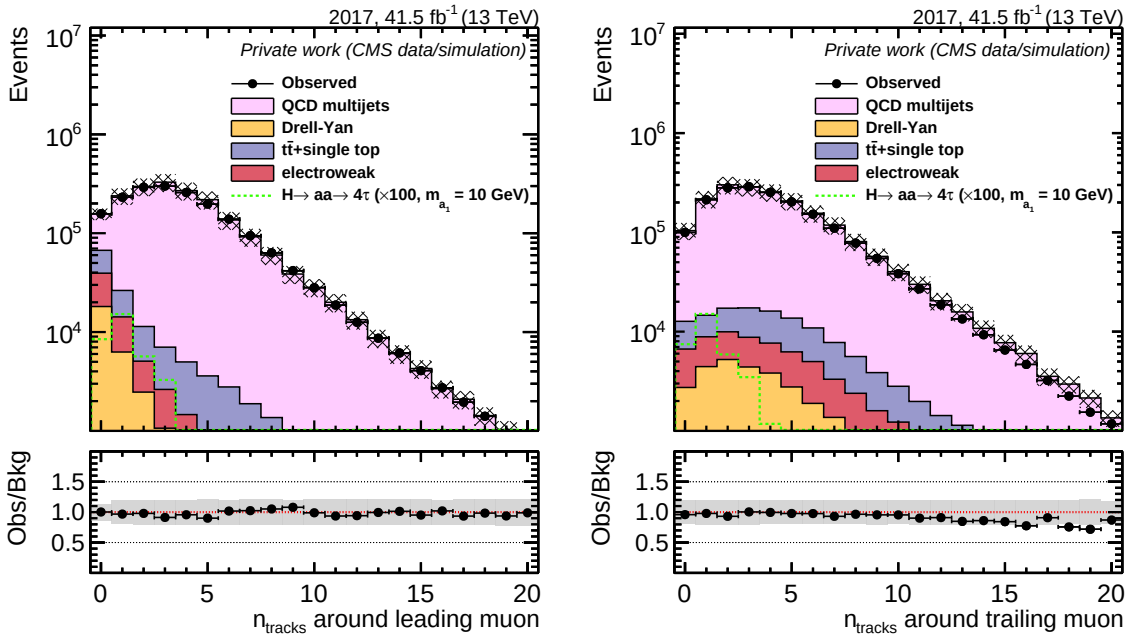

Figure 5.11: Distributions of the isolation-track multiplicity within a $\Delta R = 0.5$ around the leading (left) and subleading (right) muons. The data (represented by black dots) are compared with SM background expectations obtained from simulation (solid histograms). Background contributions from W+jets and diboson production processes are collectively aggregated under the "electroweak" category. The dashed histogram illustrates the distribution for the signal in the $\mathrm{H} \to \mathrm{a a_1} \to 4\tau$ channel for $m_{\mathrm{a_1}} = 10$ GeV. The signal normalization is computed assuming the SM Higgs production cross section and a branching ratio $\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1} \to 4\tau) = 0.05$. The signal is scaled by a factor of 100 for better visualization.

5.5.6 Selection of $\mathrm{a_1}$ Candidates

A muon-track system is accepted as an $\mathrm{a_1}$ candidate if the oppositely-charged isolation track around each muon also meets the criteria for a signal track.

5.5.7 Final Selection

Ultimately, an event is selected in the final sample if it contains two $\mathrm{a_1}$ candidates.

The total number of observed events in the signal region thus amounts to 7803, comprising 2042 events from 2016 data, 2229 from 2017 data, and 3532 from 2018 data. The number of expected background events derived from simulations closely matches the observed data. QCD multijet events constitute the dominant background in the SR, while contributions from other background sources account for approximately only 1% of the total events. However, the statistical uncertainties associated with these estimates are quite large, primarily due to the limited size of the MC sample used in the analysis.

The expected signal acceptance and yield for selected values of $m_{\mathrm{a_1}}$ are reported in Table 5.6. The acceptance is calculated relative to the total SM Higgs production cross section, incorporating contributions from the various production modes outlined in Table 5.7. The signal yields are calculated using a benchmark branching fraction of $\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau) =$

0.05 and assuming the SM H production cross sections. They are normalized to an integrated luminosity of 138 fb^{-1}. Contributions from the ggF, VBF, VH, and $\mathrm{t\bar{t}H}$ processes are summed up to estimate the total yields for the 4τ signals. Additionally, the yield for the $2\mu2\tau$ signal is calculated assuming that the partial widths of the $\mathrm{a}_1 \to \mu\mu$ and $\mathrm{a}_1 \to \tau\tau$ decays follow the relation [200]:

$$\frac{\Gamma(\mathrm{a}_1 \to \mu\mu)}{\Gamma(\mathrm{a}_1 \to \tau\tau)} = \frac{m_\mu^2}{m_\tau^2\sqrt{1-(2m_\tau/m_{\mathrm{a}_1})^2}}. \tag{5.2}$$

Using this relation between the partial widths, the ratio of branching fractions for $\mathrm{a}_1\mathrm{a}_1 \to 2\mu2\tau$ and $\mathrm{a}_1\mathrm{a}_1 \to 4\tau$ decays is calculated as

$$\frac{\mathcal{B}(\mathrm{a}_1\mathrm{a}_1 \to 2\mu2\tau)}{\mathcal{B}(\mathrm{a}_1\mathrm{a}_1 \to 4\tau)} = 2\frac{\mathcal{B}(\mathrm{a}_1 \to \mu\mu)}{\mathcal{B}(\mathrm{a}_1 \to \tau\tau)} = 2\frac{\Gamma(\mathrm{a}_1 \to \mu\mu)}{\Gamma(\mathrm{a}_1 \to \tau\tau)}. \tag{5.3}$$

The factor of 2 in Eq. 5.3 accounts for the two possible decay combinations: ${\mathrm{a}_1}^{(1)}{\mathrm{a}_1}^{(2)} \to 2\mu2\tau$ and ${\mathrm{a}_1}^{(1)}{\mathrm{a}_1}^{(2)} \to 2\tau2\mu$, both yielding final states with two muons and two tau leptons. The calculated ratio varies from approximately 0.0073 at $m_{\mathrm{a}_1} = 15\,\mathrm{GeV}$ to 0.0155 at $m_{\mathrm{a}_1} = 4\,\mathrm{GeV}$. To account for contributions from production modes other than ggF to the $2\mu2\tau$ final state, the analysis assumes identical acceptance ratios between the $2\mu2\tau$ and 4τ channels for all production mechanisms and scales the ggF acceptance accordingly.

m_{a_1} [GeV]	Acceptance [10^{-4}]		Number of events	
	4τ	$2\mu2\tau$	4τ	$2\mu2\tau$
5	3.52 ± 0.10	103.2 ± 1.2	133.9 ± 3.8	39.7 ± 0.4
8	2.55 ± 0.09	76.0 ± 1.0	97.2 ± 3.3	23.0 ± 0.3
12	1.37 ± 0.06	35.6 ± 0.7	52.1 ± 2.4	10.1 ± 0.2
15	0.32 ± 0.03	7.5 ± 0.3	12.3 ± 1.1	2.1 ± 0.1

Table 5.6: The signal acceptance and the expected number of signal events after selection in the SR. The acceptance is calculated relative to the total H production cross section, utilizing values predicted by the SM. The expected number of signal events is determined using a benchmark branching fraction of $\mathcal{B}(\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1)\mathcal{B}^2(\mathrm{a}_1 \to \tau\tau) = 0.05$, assuming SM-predicted Higgs production cross sections. The number of events is normalized to an integrated luminosity of 138 fb^{-1}. The quoted uncertainties for these predictions include only statistical uncertainties.

Process	Cross section (pb)
ggF	48.58
VBF	3.78
VH	2.26
$\mathrm{t\bar{t}H}$	0.51

Table 5.7: The standard model Higgs boson production cross sections used in the analysis for the various mechanisms [46].

5.6 Corrections to Simulations

Monte Carlo simulations are extensively used in this analysis in modeling of the signal, as discussed later in this chapter. Also, as outlined in the previous section, they are used to design and optimize the selection criteria in the signal region. Therefore accurate event description by MC simulations is of utmost importance. However, MC samples often cannot fully replicate all physical aspects of real data, necessitating corrections to mitigate the differences. These corrections address discrepancies arising from detector imperfections, challenges in modeling pileup, and limitations in the precision of theoretical calculations by MC generators. This section provides a comprehensive description of the corrections applied to the simulated samples used in this analysis[3].

5.6.1 Pileup Reweighting

Since the simulated events might not perfectly match the actual pileup conditions observed in the data collected by a detector, the weight of each simulated event is adjusted to correct for the differences between the pileup distributions in recorded data and simulation.

5.6.2 Prefiring Weights

During Run 2, an increase in the offset of the ECAL timing pulse caused the trigger primitives in the region $2.5 < |\eta| < 3.0$ to be erroneously assigned to the previous bunch crossing. Due to L1 rules forbidding events from firing in two consecutive bunch crossings, this misassignment can lead to events self-vetoing. A similar timing issue affected the bunching crossing assignments of muon candidates due to the limited temporal resolution of the muon system. To emulate this effect in simulation, event weights are adjusted [201].

5.6.3 Tracking Efficiency Corrections

Efficiencies for electron and muon tracking often vary between observed data and MC simulations. To minimize these discrepancies, centrally provided scale factors are applied to the simulated samples.

5.6.4 Muon Momentum Scale Corrections

Muon momentum measurements can be impacted by inaccuracies caused by detector anomalies, magnetic field inhomogeneities, and interaction with detector materials. Hence, the momentum scale is corrected in both data and simulation using the so-called Rochester correction factors [202].

5.6.5 Muon Identification Efficiency Corrections

This efficiency is corrected to account for the mismodeling of muon reconstruction and identification in simulation. Scale factors provided centrally for the medium working point are used.

[3]Note: The MC samples used for the offline selection, discussed previously, have already been adjusted to include the necessary corrections. These corrections are, however, detailed later in the chapter to ensure that the rationale behind each correction is understood in the context of the offline selection process.

5.6.6 Trigger Efficiency

The efficiencies for the double muon triggers employed in this analysis, listed in Table. 5.4, are calculated using the well-established tag-and-probe method. This technique involves selecting pairs of particles from a clean and well-understood process. One particle in the pair, called the "tag", is required to pass very strict selection criteria to ensure it is genuinely the particle of interest. The other particle, the "probe," is subject to much looser criteria and is used to test the efficiency of specific responses, such as the trigger selection in this context. The efficiency is determined by the ratio of the number of probes that pass the specific criteria to the total number of probes tested.

For the measurements described here, a sample of $Z \to \mu\mu$ events is utilized. As each trigger comprises two muon p_T legs, the efficiencies of each leg are calculated individually and then combined. Details of how these efficiencies are calculated for each trigger are further elaborated in the subsequent subsections. The tag-and-probe measurement is performed separately on both data and Drell–Yan MC events. The ratio of efficiencies, $\epsilon_{\text{data}}/\epsilon_{\text{MC}}$, is then applied as a weight on an event-by-event basis to correct the simulation.

5.6.6.1 Same-sign double muon triggers in 2016 and 2018

For the double muon triggers employed in 2016 and 2018, the efficiencies of each muon p_T leg are measured relative to the offline selection as a function of the muon p_T and $|\eta|$. The tag muon is chosen to be an offline muon that triggers the `HLT_IsoMu27` path[4]. It must meet additional requirements such as $p_T > 28\,\text{GeV}$, $|\eta| < 2.4$, passing a medium muon ID, strict offline isolation criteria, and impact parameters $|d_)| < 0.05\,\text{mm}$ and $|d_z| < 1\,\text{mm}$.

The probe muon is required to adhere to the nominal selection criteria used in the analysis. Additionally, the tag and probe muons must have opposite charges and a minimum angular separation of $\Delta R > 1.0$.

Given that the `HLT_IsoMu27` trigger imposes a more stringent p_T requirement than either leg of the double muon trigger, the tag muon will inherently satisfy the conditions of both legs of the double muon trigger. Therefore, the efficiency of each leg can be determined by evaluating how frequently the probe muon triggers the respective leg. The efficiency results for both legs of the same-sign double muon triggers in 2018 are presented in Figures 5.12 and 5.13. The results for 2016 are available in Appendix D.

The same-sign efficiency of the trigger is calculated by subtracting the online charge misidentification rate from unity. The charge misidentification rate is determined as the fraction of $Z \to \mu^+\mu^-$ candidates (which contain two opposite-sign muons) that pass the same-sign filter.

5.6.6.2 Double muon trigger in 2017

The trigger `HLT_Mu17_TrkVVL_Mu8_TrkVVL_DZ_Mass8` employed in the analysis of the 2017 data sample sets p_T thresholds for both the leading and subleading muons and requires loose tracker isolation ($I_\mu < 0.4$) for each, as elaborated in Sec 5.4. The efficiency of each muon

[4]This trigger selects events that contain at least one muon with $p_T > 27\,\text{GeV}$ and imposes a tight online isolation criterion of $I_\mu < 0.1$ on it.

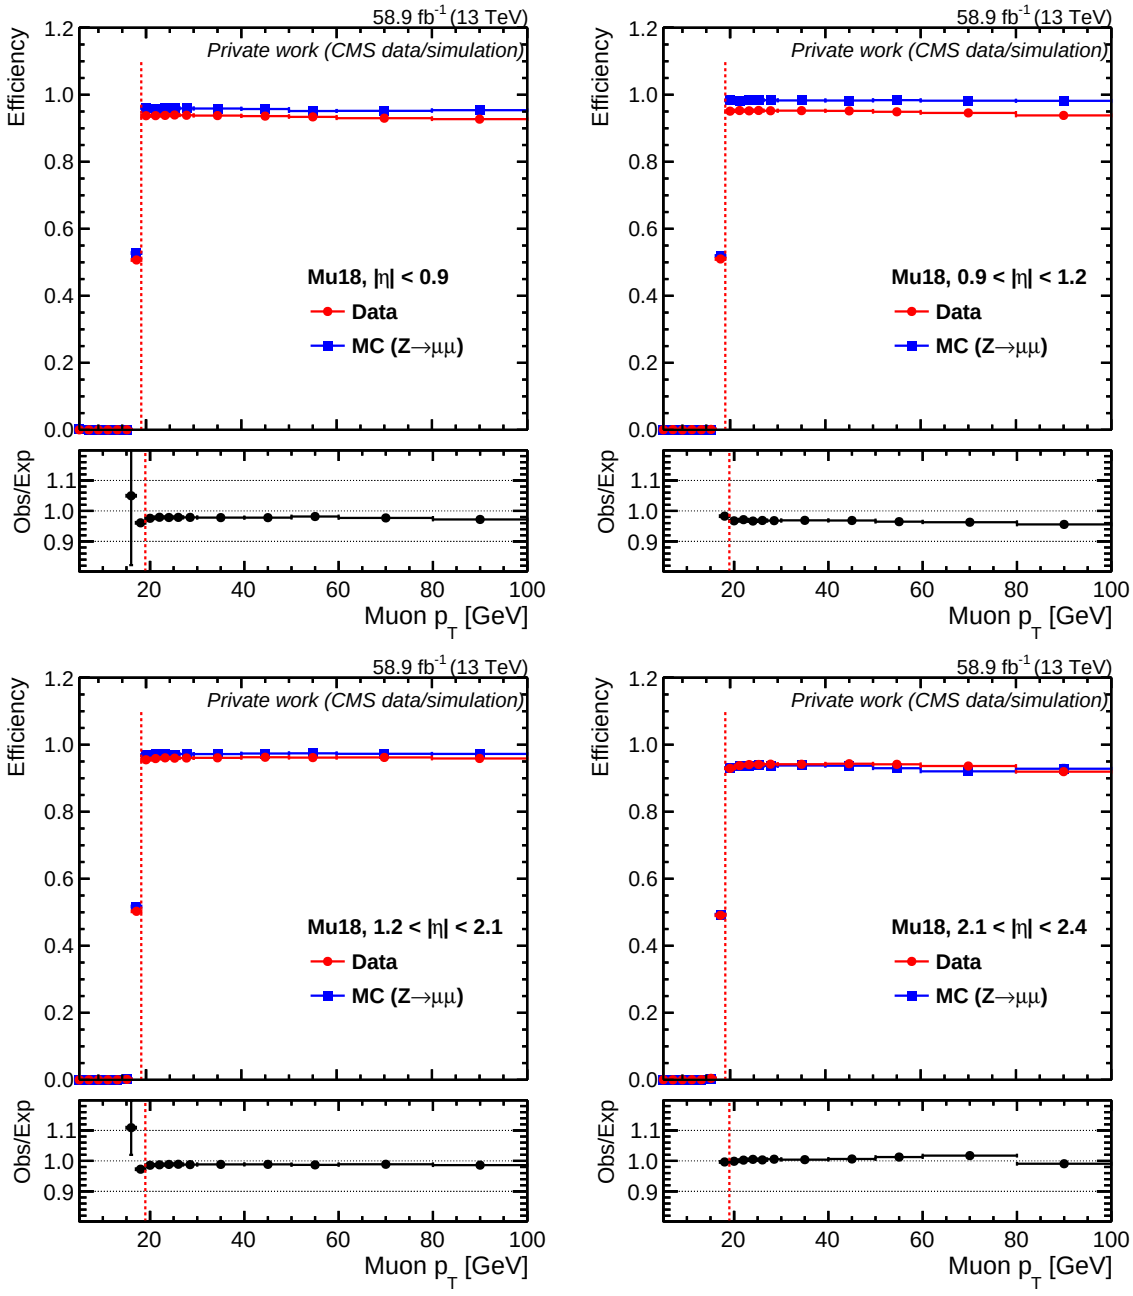

Figure 5.12: Efficiency of the `Mu18` leg of the `HLT_Mu18_Mu9_SameSign` trigger as a function of the muon transverse momentum for four different eta bins: $|\eta| < 0.9$ (top left), $0.9 < |\eta| < 1.2$ (top right), $1.2 < |\eta| < 2.1$ (bottom left), and $|\eta| > 2.1$ (bottom right). The circles present measurements in data, and the squares are measurements in MC.

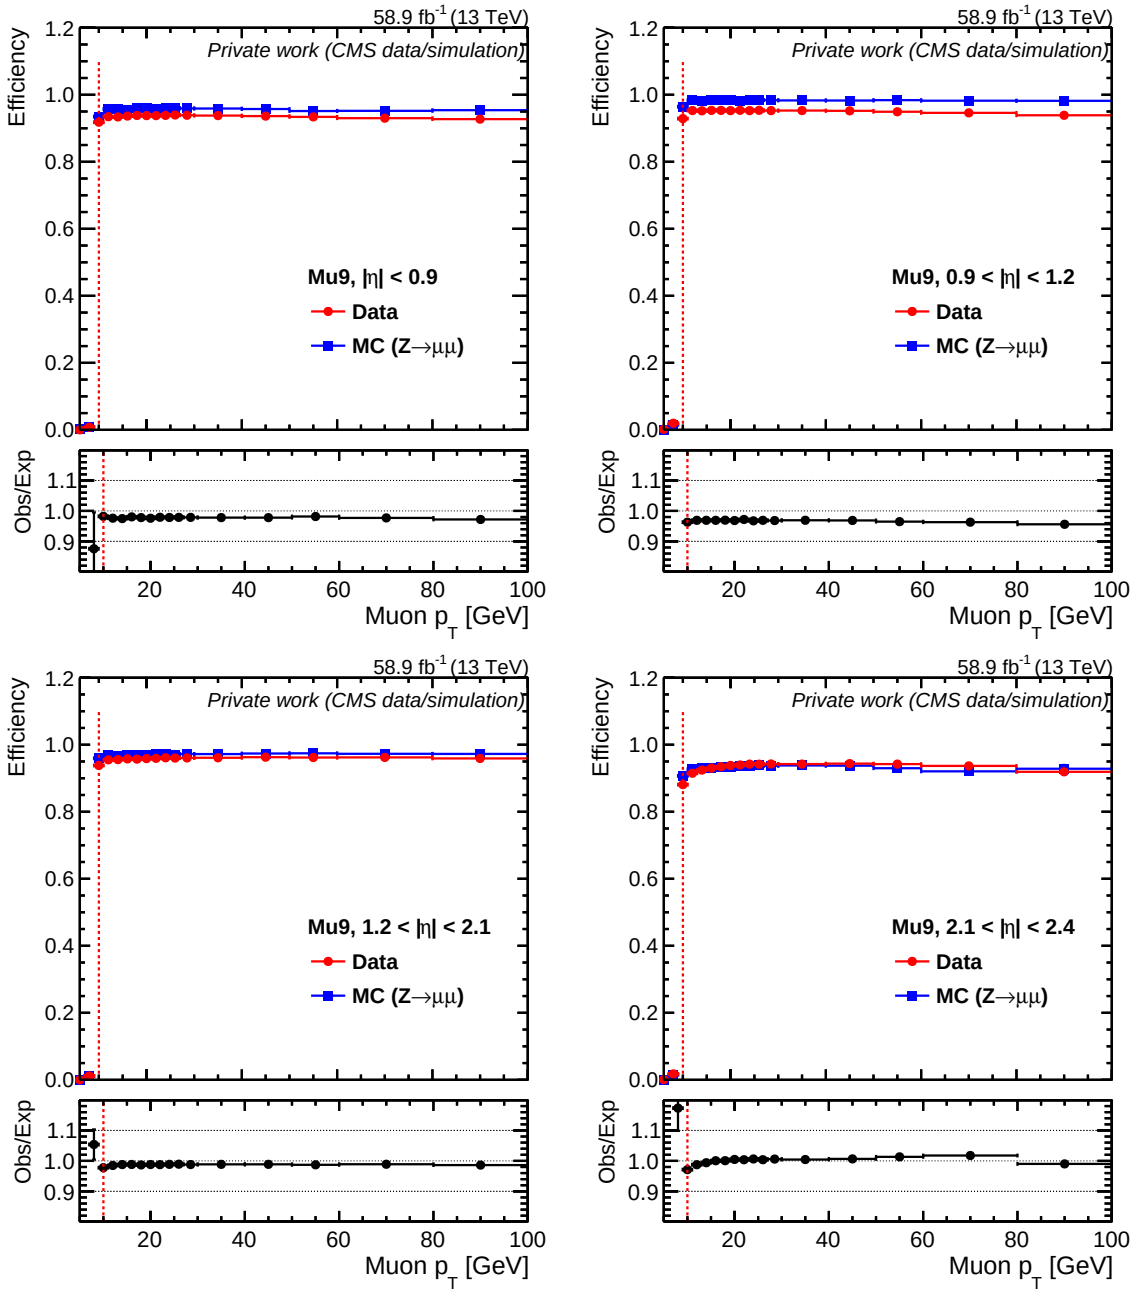

Figure 5.13: Efficiency of the `Mu9` leg of the trigger `HLT_Mu18_Mu9_SameSign` as a function of the muon transverse momentum for four different eta bins: $|\eta| < 0.9$ (upper left plot), $0.9 < |\eta| < 1.2$ (upper right plot), $1.2 < |\eta| < 2.1$ (lower left plot), and $|\eta| > 2.1$ (lower right plot). The circles present measurements in data, and the squares are measurements in MC.

leg (`Mu17_TrkVVL` and `Mu8_TrkVVL`) is thus calculated as a product of two components: the efficiency of the corresponding p_{T} filters and the efficiency of the isolation requirement.

The efficiency of the p_{T} filters, measured as a function of the muon's p_{T} and η, is determined using the same methodology as for the same-sign double muon triggers. The plots displaying these measurements are provided in Appendix D.

The isolation efficiency for each leg is calculated from the subset of probe muons that have passed the corresponding p_{T} filter and also meet the online isolation requirement. Since the tag muon is required to pass a single muon trigger with a stricter online isolation criterion ($I_\mu < 0.1$), the isolation efficiency for the double muon trigger depends solely on whether the probe muon satisfies the looser online isolation criteria of the double muon trigger. This efficiency is measured as a function of the isolation variable, I_μ, derived from tracks reconstructed offline. The results are illustrated in Figure 5.14.

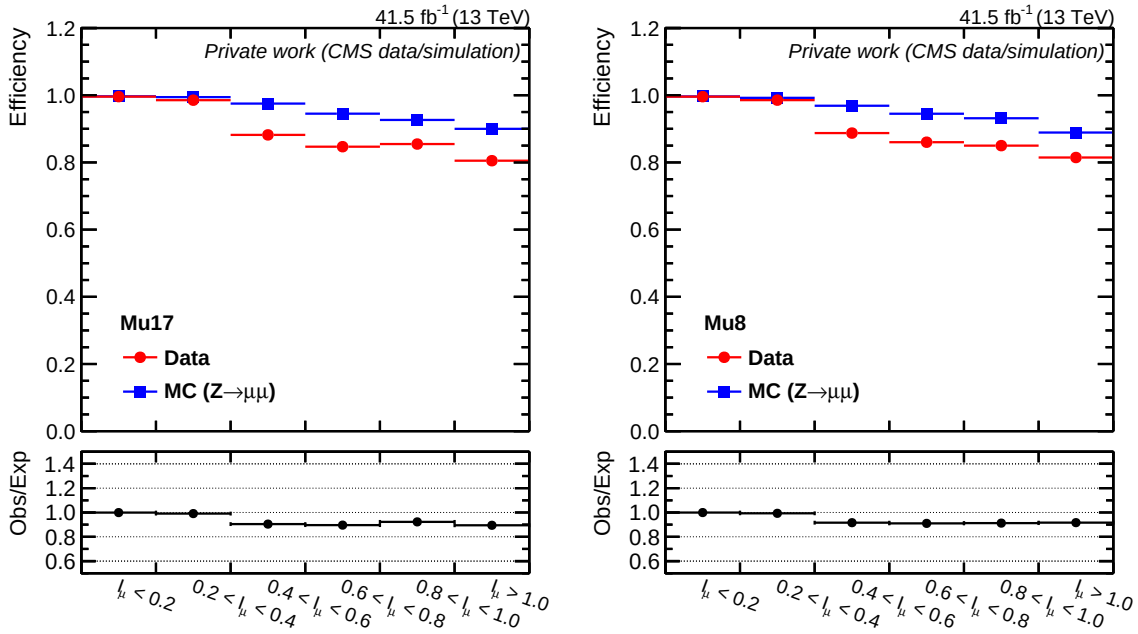

Figure 5.14: Isolation efficiency for the double muon trigger in the 2017 data and simulated samples presented as a function of the muon isolation variable, I_μ, calculated using offline tracks. The efficiencies for the `Mu17` leg and the `Mu8` leg are displayed in the left and right panels, respectively. The lower panels of each plot depict the corresponding data-MC scale factors.

5.6.7 Track Isolation and Identification Efficiency

The isolation criteria for muon-track pairs can be influenced by the presence of additional charged particles from pileup (PU) interactions or underlying events, potentially leading to discrepancies between MC simulations and actual data. Such differences could, in turn, also affect the efficiency of selecting one-prong tau candidates. To assess these potential inconsistencies, a detailed study was conducted using a sample of $\mathrm{Z} \to \tau\tau$ events, especially focusing on the $\mathrm{Z} \to \tau_\mu \tau_{\text{1-prong}}$ decay mode. This study aimed to closely replicate the selection criteria for the one-prong tau candidates used in the analysis. The exact strategy and procedures

followed are provided in Appendix F. Figure 5.15 shows the correction scale factors thus derived for different track p_{T} ranges for 2016. The scale factors derived for the other two years are given in Appendix F. As no significant p_{T} dependence is observed, a constant scale factor (indicated by the dark blue line) was applied across all simulated events. To account for potential variations, p_{T}-dependent uncertainties were modeled using a linear fit and incorporated into the final corrections.

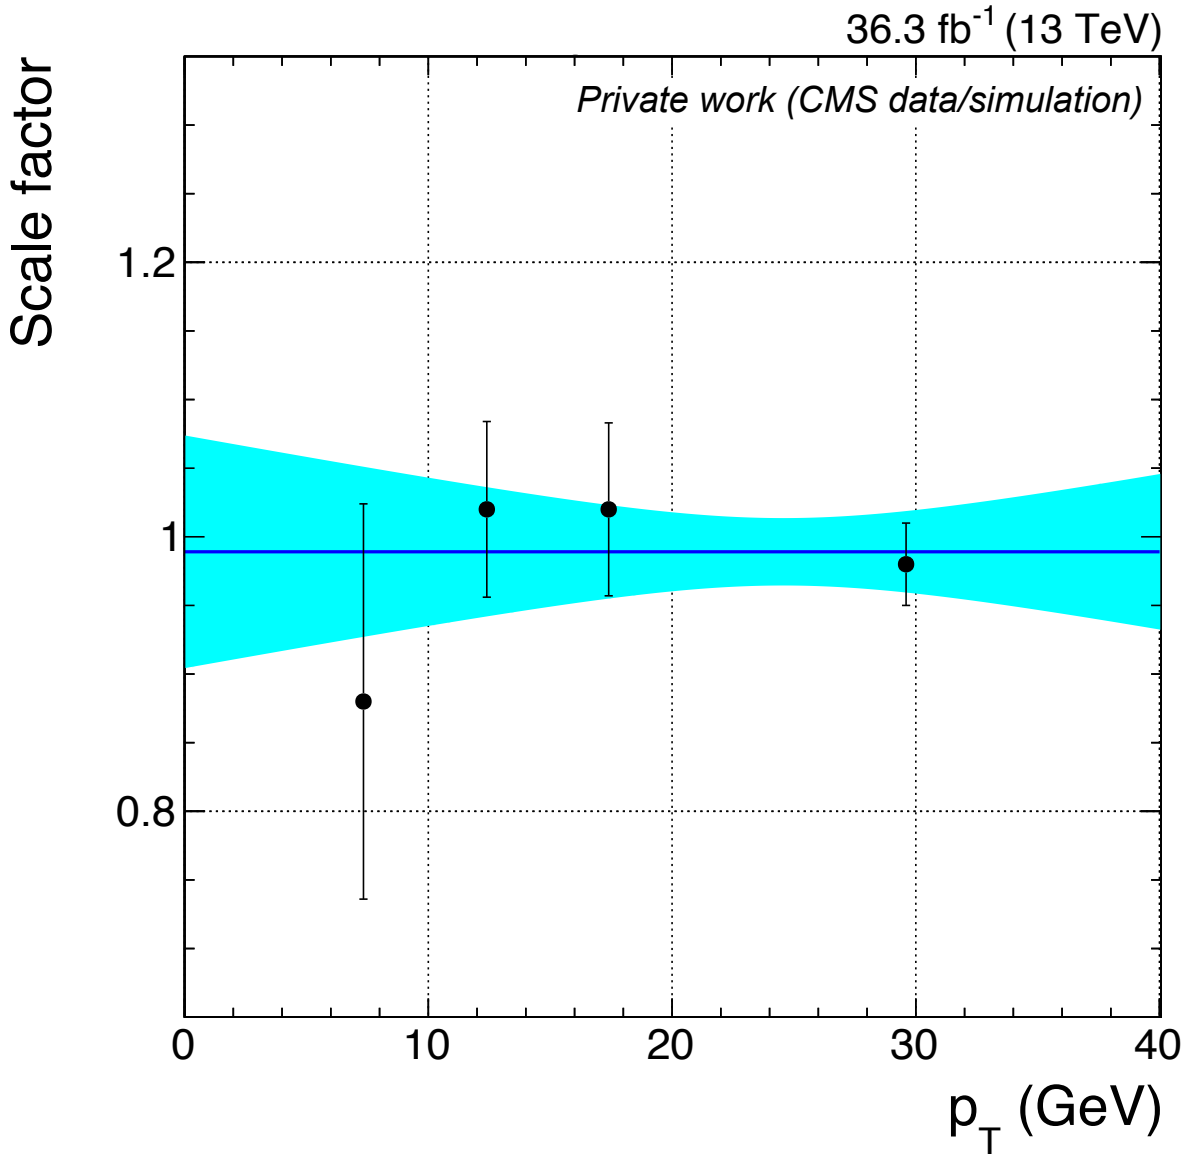

Figure 5.15: The scale factor for the track isolation and identification efficiency as a function of the track p_{T} for 2016. As no dependence of the scale factors is observed with respect to the track p_{T}, the scale factors are fitted with a constant function (dark blue line). The blue band represents the uncertainty associated with the scale factor at each p_T value, modeled using a linear fit.

5.6.8 Higgs-p_T Reweighting

As mentioned in Sec. 5.3, the simulated signal samples are generated at leading order in QCD. The signal acceptance, however, is found to be moderately affected by the p_T of the Higgs, as variations in this parameter can result in the a_1 bosons being more or less boosted. This, in turn, affects the angular separation ΔR of the muon-track pairs. To account for this, each event in the simulated signal samples is reweighted with a k-factor to match the higher-order predictions for the p_T distribution of the H boson across all production modes. For example, Fig. 5.16 presents the NNLO p_T distribution of the Higgs boson for the gluon-gluon fusion production mode obtained with the HQT (v.2.0) program [203, 204], alongside the LO spectra from PYTHIA (v.8.212). Also shown are the k-factors obtained by dividing the two. The higher-order predictions for the Higgs p_T distributions and k-factors for the other production modes are presented in Appendix E.

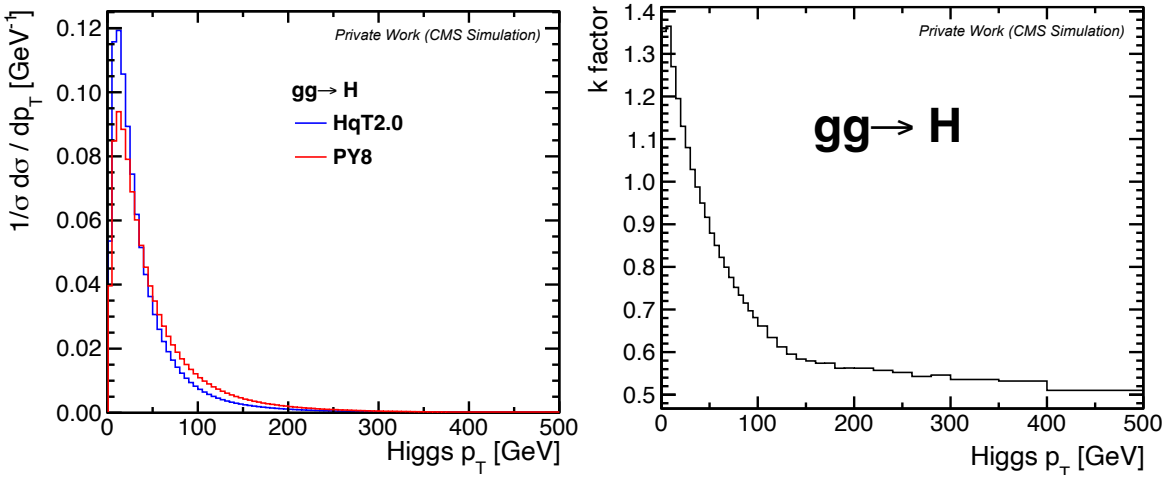

Figure 5.16: The distribution of the Higgs p_T as predicted at LO by PYTHIA (v.8.212) and at NNLO by HQT (v.2.0) in the gluon-gluon fusion process (left) and corresponding k-factor (right).

5.6.9 Jet Energy Corrections and b-jet Identification Efficiency

Imposing a b-jet veto in the analysis necessitates precise corrections to both the jet energy scale (JES) and b tagging efficiency to identify and exclude b-jets accurately. The energy of jets in both data and Monte Carlo simulations is corrected by applying JES uncertainties as provided centrally by CMS. Furthermore, b tagging efficiency is carefully corrected for in simulation by reweighting the events in accordance with the recommendations set forth by CMS [205].

5.7 Signal Extraction Stategy

In order to discriminate between the signal and the background, the two-dimensional (2D) distribution of the invariant masses of the muon-track systems, which constitute the two a_1 candidates, is used. Due to the presence of several neutrinos in both a_1 decay legs of the 4τ channel, the full reconstruction of the mass of a_1 is impossible. This precludes the usage of

analytic parametric models for describing the resonant mass distributions and necessitates the adoption of a binned 2D distribution approach instead.

The binning of the 2D distribution used in the analysis is illustrated in Fig. 5.18. For simplicity, this distribution is filled with pairs of muon-track invariant masses (m_1, m_2), ordered such that $m_2 > m_1$. Consequently, only the bins (i, j) where $j \geq i$ are filled, creating in total $6(6+1)/2 = 21$ independent bins. The bin boundaries along each axis are set at (0.0, 1.0, 2.0, 3.0, 4.0, 5.2) GeV. Specifically, bins $(i, 6)$ with $i = 1, ..., 5$ capture all events where $m_2 > 5.2\,\text{GeV}$, and bin $(6, 6)$ contains all events where both $m_{1,2} > 5.2\,\text{GeV}$. Muon-track invariant masses below 4 GeV (the lowest a_1 mass hypothesis) are considered to account for the dispersion caused in the invariant mass distributions by the presence of undetected neutrinos in the a_1 decay legs, as illustrated in Fig. 5.17. The binning is further optimized to ensure high and uniform signal sensitivities across all pseudoscalar mass hypotheses. It also maintains reasonable statistics in each bin of the 2D (m_1, m_2) distribution in the control regions used to construct the background model, described in Sec. 5.9.

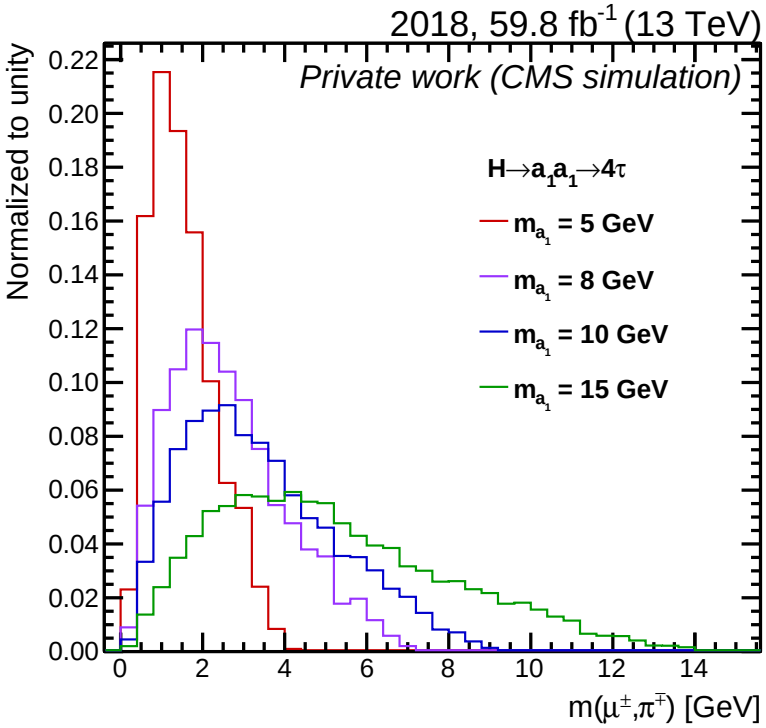

Figure 5.17: Distributions of the invariant masses of generator-level muon-pion pairs in the $H \to a_1 a_1 \to 4\tau$ process, for different a_1 masses.

The modeling of the binned 2D (m_1, m_2) distributions is performed for the signal processes using simulated samples, while the background modeling employs a dedicated data-driven approach. Further details on these procedures are explained in Sections 5.8 and 5.9.

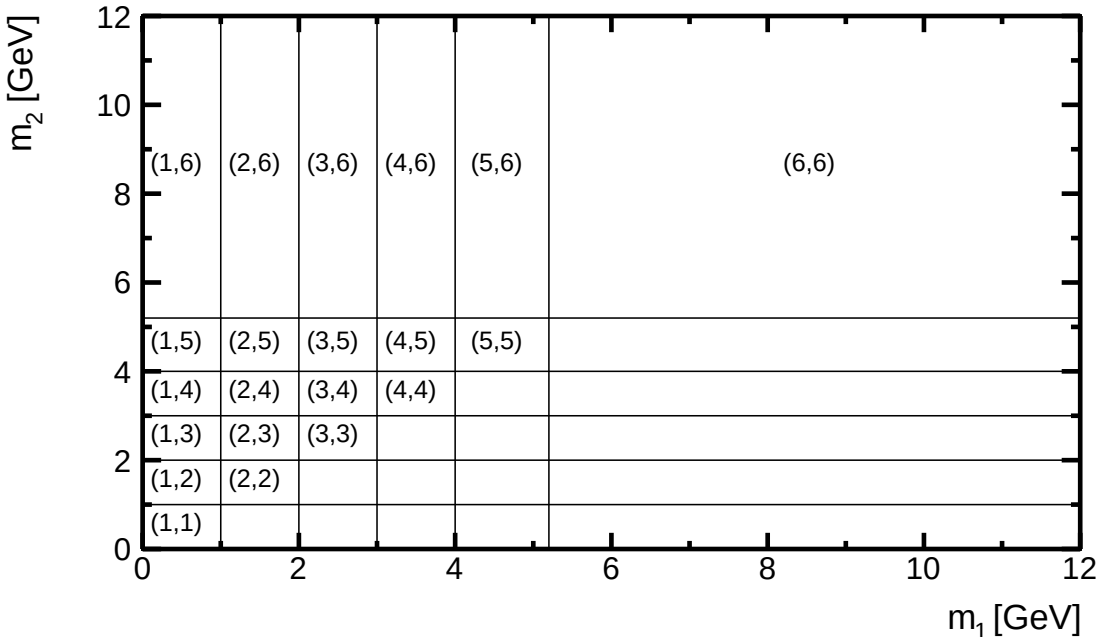

Figure 5.18: Binning of the 2D (m_1, m_2) distribution.

5.8 Signal Modeling

The 2D invariant mass distributions for the signal are constructed using simulated samples of the $\mathrm{H} \to \mathrm{a_1a_1} \to 4\tau$ and $\mathrm{H} \to \mathrm{a_1a_1} \to 2\mu 2\tau$ decays. The relative contributions from the different processes corresponding to the different Higgs boson production modes are determined based on their respective cross sections as predicted by the SM. The contribution from the $\mathrm{H} \to \mathrm{a_1a_1} \to 2\mu 2\tau$ decay is computed under the assumption that the partial widths of the $\mathrm{a_1} \to \tau\tau$ and $\mathrm{a_1} \to \mu\mu$ decays satisfy Eq. 5.2.

Figures 5.19–5.21 display the normalized 2D invariant mass distributions for the $\mathrm{H} \to \mathrm{a_1a_1} \to 4\tau$ and $\mathrm{H} \to \mathrm{a_1a_1} \to 2\mu 2\tau$ signal samples, presented as a one-dimensional array for the mass hypotheses of $m_{\mathrm{a_1}} = 5$, 8, 12 and 15 GeV. The binning here follows the notation of Fig 5.18. The signal distributions are normalized to the SM predictions for the Higgs boson production rate, with an assumed branching fraction of $\mathcal{B}(\mathrm{H} \to \mathrm{a_1a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau) = 0.05$. As evident from the figure, the shapes of the two processes differ considerably. The invariant mass distribution of the muon-track system in the $\mathrm{a_1} \to \mu\mu$ decay channel peaks at the nominal $\mathrm{a_1}$ boson mass. In contrast, the reconstructed mass of the muon-track system in the $\mathrm{a_1} \to \tau\tau$ decay is generally lower and has a larger dispersion due to the presence of undetected neutrinos.

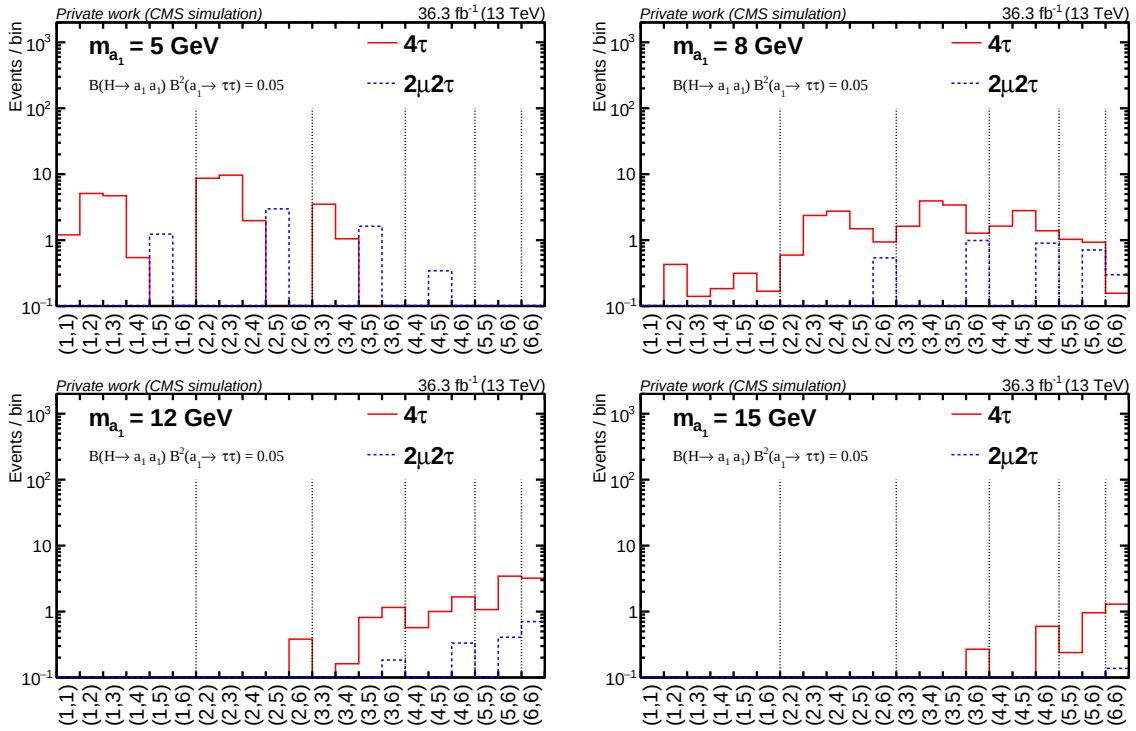

Figure 5.19: The signal $f_{2D}(i,j)$ templates for mass hypothesis $m_{a_1} = 5\,\text{GeV}$ (upper left), 8 GeV (upper right), 12 GeV (lower left) and 15 GeV (lower right), shown for the year 2016. The plots feature the $H \to a_1a_1 \to 2\mu 2\tau$ contributions in blue and the $H \to a_1a_1 \to 4\tau$ contributions in red. These distributions are normalized based on the SM H production cross section with an assumed branching fraction $\mathcal{B}(H \to a_1a_1)\mathcal{B}^2(a_1 \to \tau\tau) = 0.05$. The binning follows the notation of Fig. 5.18.

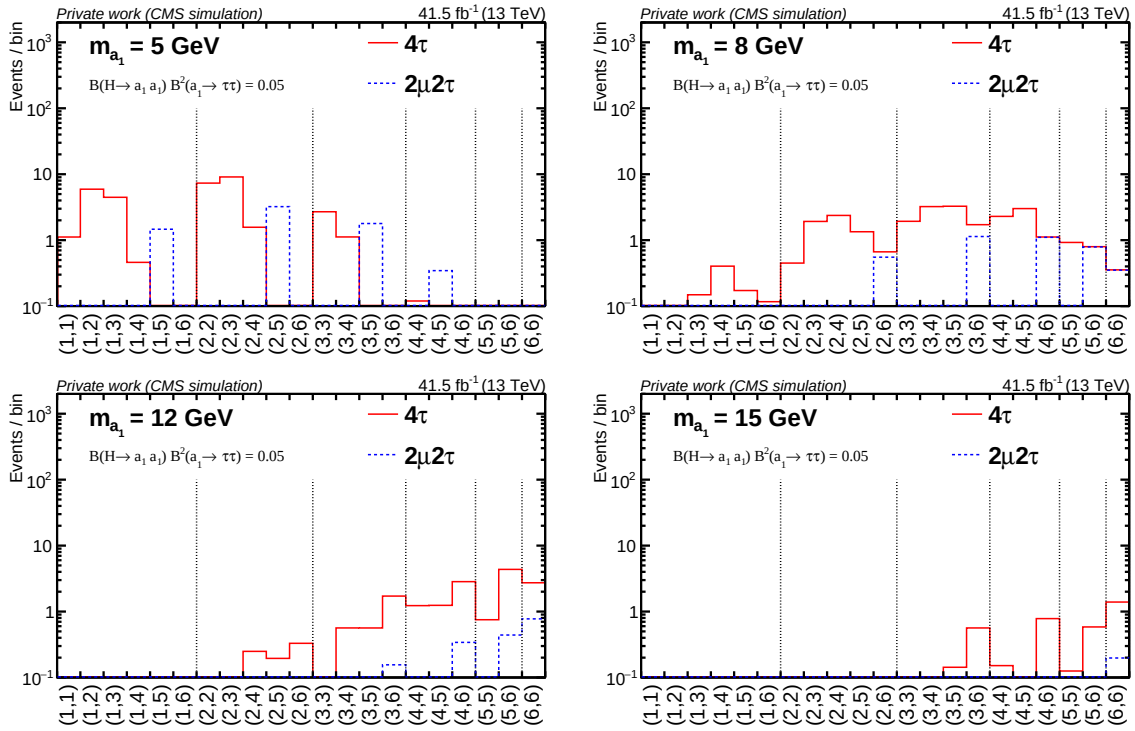

Figure 5.20: The signal $f_{2D}(i,j)$ templates for mass hypothesis $m_{a_1} = 5\,\text{GeV}$ (upper left), 8 GeV (upper right), 12 GeV (lower left) and 15 GeV (lower right), shown for the year 2017. The plots feature the $H \to a_1a_1 \to 2\mu 2\tau$ contributions in blue and the $H \to a_1a_1 \to 4\tau$ contributions in red. These distributions are normalized based on the SM H production cross section with an assumed branching fraction $\mathcal{B}(H \to a_1a_1)\mathcal{B}^2(a_1 \to \tau\tau) = 0.05$. The binning follows the notation of Fig. 5.18.

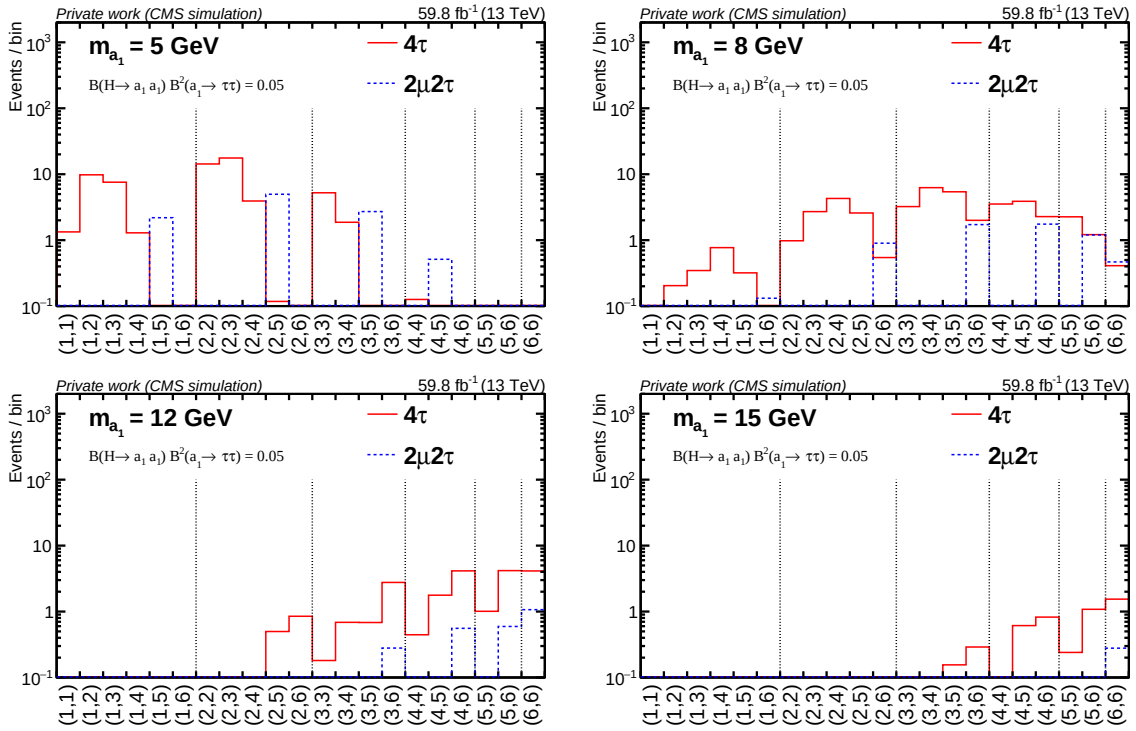

Figure 5.21: The signal $f_{2D}(i,j)$ templates for mass hypothesis $m_{a_1} = 5\,\text{GeV}$ (upper left), 8 GeV (upper right), 12 GeV (lower left) and 15 GeV (lower right), shown for the year 2018. The plot features the $H \to a_1a_1 \to 2\mu 2\tau$ contributions in blue and the $H \to a_1a_1 \to 4\tau$ contributions in red. These distributions are normalized based on the standard model Higgs production cross section with an assumed branching fraction $\mathcal{B}(H \to a_1a_1)\mathcal{B}^2(a_1 \to \tau\tau) = 0.05$. The binning follows the notation of Fig. 5.18.

5.9 Background Modeling

An accurate description of backgrounds is critical for distinguishing authentic signals from mere background fluctuations in experimental data. As mentioned in Section 5.5.7, QCD multijets form the dominant source of background in this analysis, with minor contributions coming from processes such as $t\bar{t}$, Z+jets, W+jets, and diboson production. However, QCD processes are known to be extremely challenging to simulate with precision. Also, the small size of the QCD MC samples makes them impractical for estimating background. Therefore, a data-driven approach is used to model the total background in the SR, including both QCD and non-QCD processes.

To model the shape of the 2D (m_1, m_2) distribution of the background in the SR, a binned template is constructed as:

$$f_{\mathrm{2D}}(i,j) = C(i,j)(\, f_{\mathrm{1D}}(i) f_{\mathrm{1D}}(j))^{\mathrm{sym}}, \tag{5.4}$$

where

- $f_{\mathrm{2D}}(i,j)$ refers to the content of the bin (i,j) in the normalized 2D invariant mass distribution of the muon-track systems;
- $f_{\mathrm{1D}}(i)$ is the content of bin i in the normalized one-dimensional (1D) distribution of the muon-track invariant mass;
- $C(i,j)$ is a symmetric matrix, accounting for possible correlation between m_1 and m_2. The condition $C(i,j) = 1$ for all bins (i,j) would indicate an absence of correlation between m_1 and m_2. The elements of the matrix $C(i,j)$ are referred to as "correlation factors" henceforth.

The term 'sym' in Eq. 5.4 represents a symmetrization operation applied to the product of the one-dimensional distributions $f_{\mathrm{1D}}(i)$ and $f_{\mathrm{1D}}(j)$. This operation is defined as follows:

$$(f_{\mathrm{1D}}(i) f_{\mathrm{1D}}(j))^{\mathrm{sym}} = \begin{cases} 2 f_{\mathrm{1D}}(i) f_{\mathrm{1D}}(j), & \text{if } i \neq j, \\ f_{\mathrm{1D}}(i) f_{\mathrm{1D}}(i), & \text{if } i = j. \end{cases} \tag{5.5}$$

The contents of the nondiagonal bins (i,j) and (j,i) in the Cartesian product $f_{\mathrm{1D}}(i) f_{\mathrm{1D}}(j)$ are summed up to ensure that all events are accurately accounted for in the 2D (m_1, m_2) distribution while maintaining the proper ordering of the muon-track invariant masses.

Additionally, it is important to highlight that the background normalization is left unconstrained prior to the extraction of the signal.

In order to derive and validate the modeling of $f_{\mathrm{1D}}(i)$ and $C(i,j)$, multiple control regions (CRs) are defined based on varying isolation criteria applied to one or both muon-track pairs. These isolation criteria are characterized by the number of tracks within a cone of $\Delta R = 0.5$ around the muon momentum direction. According to these criteria, the two muons are categorized as the "first muon" and "second muon". A summary of the CRs used to derive and validate the modeling of the background shape, describing the specifications for the first

Control region	First μ	Second μ	Purpose
N_{23}	$N_{\text{sig}} = 1$, $N_{\text{iso}} = 1$	$N_{\text{iso}} = 2, 3$	Determination of $f_{\text{1D}}(i)$
$N_{\text{iso},2} = 1$	$N_{\text{sig}} \geq 1$, $N_{\text{iso}} > 1$	$N_{\text{sig}} = 1$, $N_{\text{iso}} = 1$	Validation and systematic uncertainty estimate of $f_{\text{1D}}(i)$
$N_{\text{iso},2} = 2, 3$	$N_{\text{sig}} \geq 1$, $N_{\text{iso}} > 1$	$N_{\text{iso}} = 2, 3$	Validation and systematic uncertainty estimate of $f_{\text{1D}}(i)$
Loose-Iso	$N_{\text{sig}} = 1, N_{\text{iso}} = 3, 4$	$N_{\text{sig}} = 1, N_{\text{iso}} = 3, 4$	Determination of $C(i, j)$
Signal region	$N_{\text{sig}} = 1$, $N_{\text{iso}} = 1$	$N_{\text{sig}} = 1$, $N_{\text{iso}} = 1$	Signal extraction

Table 5.8: Control regions used to construct and validate the background model. The symbols N_{iso} and N_{sig} denote the number of isolation and signal tracks, respectively, within a cone of $\Delta R = 0.5$ around the muon momentum direction. In cases where N_{sig} is not mentioned, there is no explicit requirement on the number of signal tracks. The last row defines the SR. The final row includes the signal region to highlight its difference from the control regions.

and second muon, is presented in Table 5.8. These CRs are carefully defined so that they have minimal signal contamination.

5.9.1 Modeling of $f_{\text{1D}}(i)$

The $f_{\text{1D}}(i)$ distribution is modeled using the N_{23} CR. This region includes events that pass the SS dimuon selection criteria and contain only one a_1 candidate formed by an isolated signal track and a muon (referred to as the first muon). The second muon in the event is required to be accompanied by either two or three nearby isolation tracks. Simulation studies indicate that the N_{23} CR is enriched in QCD events, with non-QCD backgrounds contributing less than 5% of the events. Moreover, the signal contamination remains below 3% across all bins of the $f_{\text{1D}}(i)$ distribution.

The invariant mass of the first muon and its associated track is utilized to construct the $f_{\text{1D}}(i)$ distribution. This modeling approach assumes that the kinematic distributions of the muon-track system forming the a_1 candidate are minimally affected by the isolation criteria applied to the second muon. As a result, the $f_{\text{1D}}(i)$ distribution in the N_{23} CR is expected to closely resemble that in the SR.

A direct test of this assumption with MC samples provides inconclusive results due to their limited statistics. Therefore, the hypothesis is verified using additional control regions in data labeled $N_{\text{iso},2} = 1$ and $N_{\text{iso},2} = 2, 3$. Events are selected in these CRs if the first muon has more than one isolation track ($N_{\text{iso}} > 1$), with at least one of these tracks meeting the criteria for signal tracks. As for the second muon, the $N_{\text{iso},2} = 1$ control region requires exactly one nearby signal track, similar to the SR. In contrast, the $N_{\text{iso},2} = 2, 3$ CR includes events where the second muon is accompanied by two or three isolation tracks, mirroring the isolation conditions in the $N_{\text{iso},2} = 2, 3$ region. As more than one of the tracks around the first muon in these two control regions may qualify as a signal track, two scenarios are considered: using either the signal track with the lowest p_{T} ("softest") or the highest p_{T} ("hardest") to compute the muon-track invariant mass. If only one signal track is found around the first muon, it serves as both the softest and hardest track. The invariant mass distributions of the first muon and its softest or hardest accompanying track are then compared between the

two isolation scenarios of the second muon.

The results of this study obtained with 2018 data are shown in Fig. 5.22. For every analyzed year, the invariant mass distributions for both cases (using the softest and hardest track) differ in each bin by less than 8%. This suggests that the invariant mass of the muon-track system forming an a_1 candidate is not highly sensitive to the isolation requirement on the second muon. To address any systematic effects on the modeling of the $f_{1D}(i)$ distribution in the N_{23} region, the observed differences are treated as an uncertainty in the normalized $f_{1D}(i)$ template.

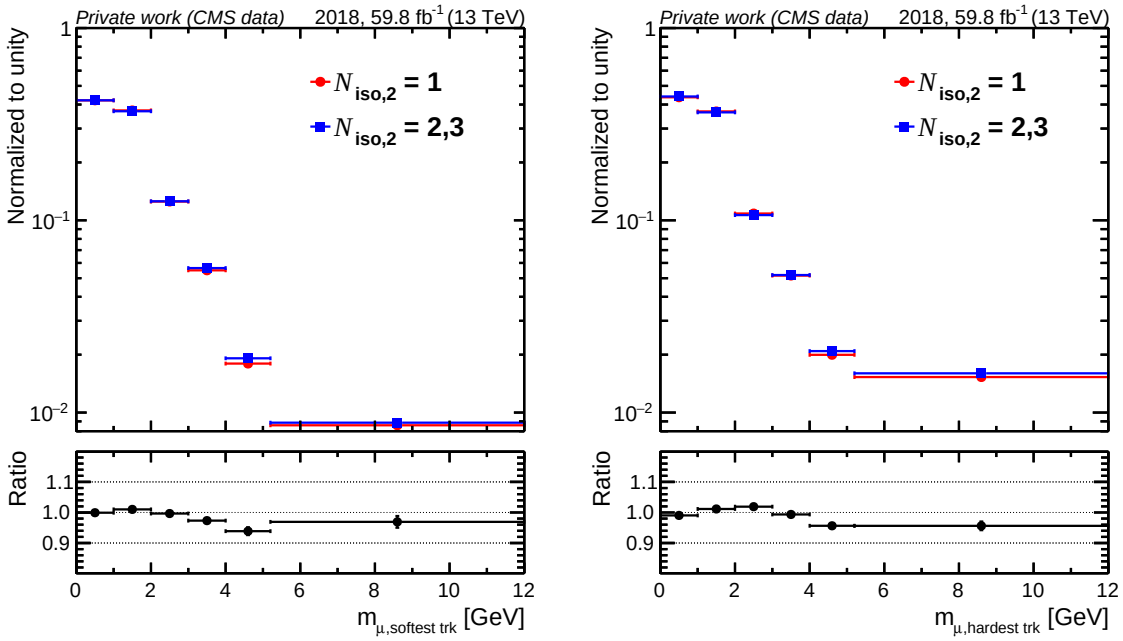

Figure 5.22: The observed invariant mass distribution, normalized to unity, of the first muon and the softest (left) or hardest (right) accompanying signal track for 2018. Two different isolation requirements are imposed on the second muon: one isolation track ($N_{\text{iso},2} = 1$) or two to three isolation tracks ($N_{\text{iso},2} = 2, 3$).

The normalized invariant mass distributions of the muon-track system for the background model, thus derived from the N_{23} CR, are presented in Figure 5.23, alongside the data selected in the SR. Each event contributes two entries, corresponding to the two muon-track systems in each event that pass the selection. The data and estimated background show good agreement within the bounds of statistical uncertainties. These distributions are also compared to the signal distributions obtained from simulation for four representative mass hypotheses, $m_{a_1} = 5$, 8, 12, and 15 GeV. It is noted that the invariant mass of the muon-track system demonstrates higher discrimination power between the background and the signal at higher m_{a_1}. For lower masses, however, the signal shape becomes more similar to the background, thereby diminishing its discriminative capability.

5.9.2 Modeling of $C(i, j)$

The correlation factor $C(i, j)$ is calculated using a different control region labeled as Loose-Iso. This CR does not overlap with the signal region and contains events featuring two same-sign

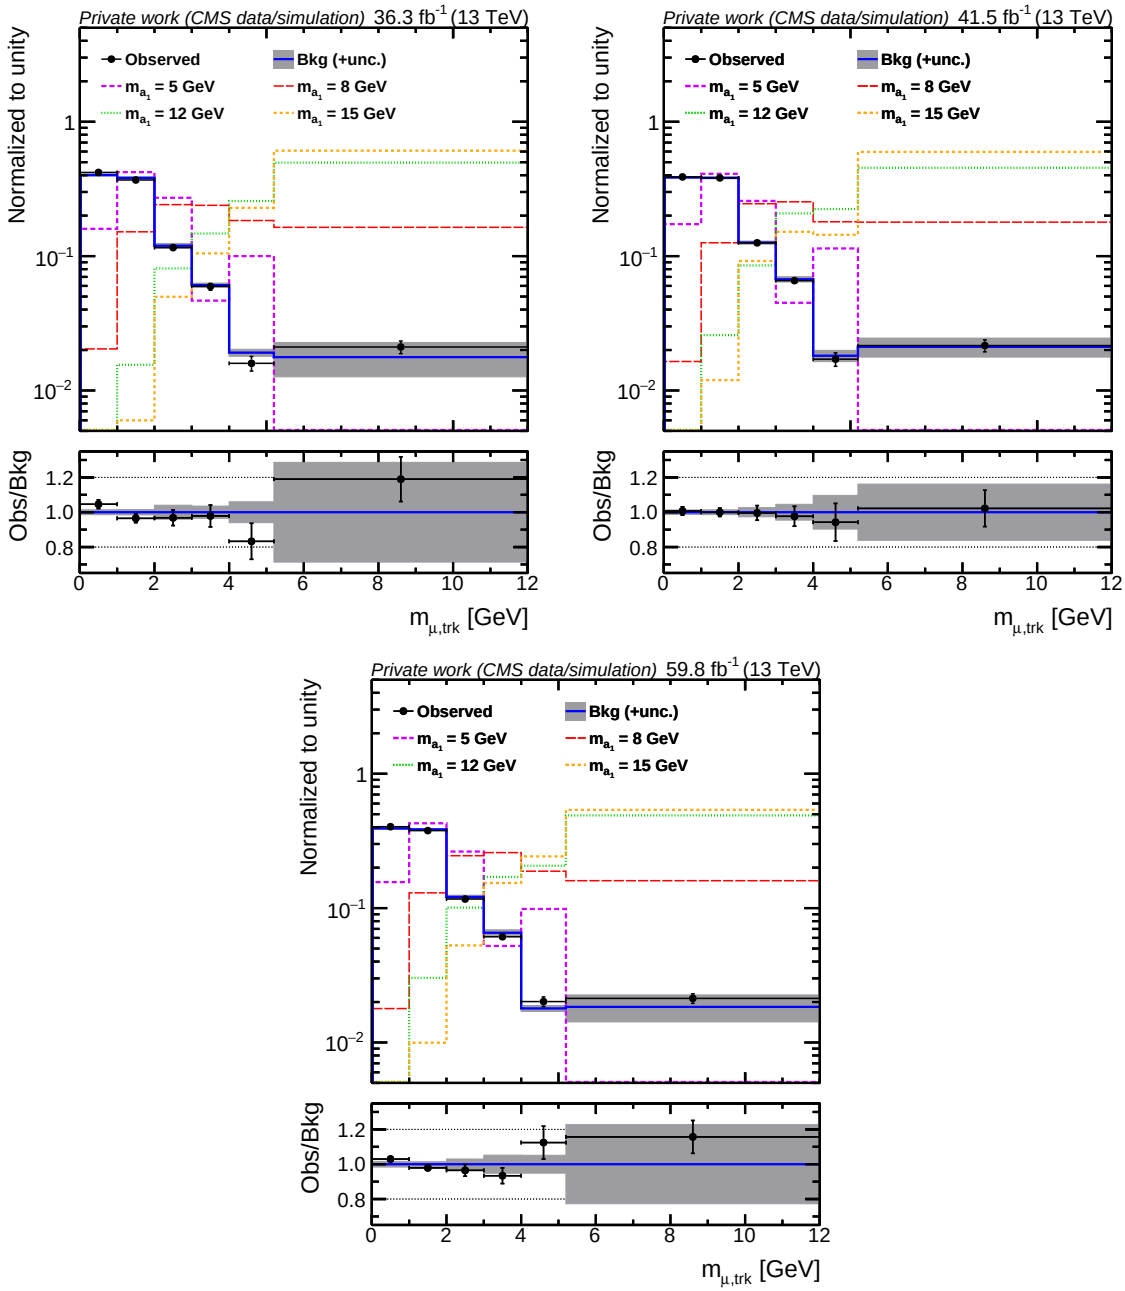

Figure 5.23: Normalized invariant mass distribution of the muon-track system for events passing the signal selection, shown for 2016 (top left), 2017 (top right), and 2018 (bottom). Observed data are represented as black points with error bars, while the background model, derived from the N_{23} control region, is illustrated in blue. Additionally, dashed histograms show the normalized distributions from signal simulations for four mass hypotheses, m_{a_1} = 5, 8, 12, and 15 GeV. Signal distributions include both the $2\mu2\tau$ and 4τ contributions. Each event contributes two entries, corresponding to the two muon-track systems in each event that pass the selection. The lower panel shows the ratio of observed to expected background events per bin. The grey shaded area represents the uncertainty of the background model.

muons that fulfill the selection criteria detailed in Sec. 5.5.3. Each muon is required to be associated with three or four nearby tracks, among which one must be classified as a signal track and the others as isolation tracks. Simulation studies reveal that QCD multijet events predominantly populate this control region, comprising approximately 99% of the events. The signal contamination in this region remains below 1% in most bins, with a modest increase to about 8% in the highest mass bin.

Events from this control region are utilized to construct the normalized 2D distributions, $f_{\mathrm{2D}}(i,j)$. They are also used to calculate the 1D normalized distribution, $f_{\mathrm{1D}}(i)$, where each event contributes two entries. The correlation factors $C(i,j)$ in the Loose-Iso region are subsequently derived using Eq. 5.4 as follows:

$$C(i,j) = \frac{f_{\mathrm{2D}}(i,j)}{(f_{\mathrm{1D}}(i) f_{\mathrm{1D}}(j))^{\mathrm{sym}}}. \tag{5.6}$$

Figures 5.24, 5.25, and 5.26 illustrate the correlation factors $C(i,j)^{\mathrm{CR}}_{\mathrm{data}}$ obtained from the Loose-Iso region for 2016, 2017, and 2018 data, respectively. Each $C(i,j)$ value is also accompanied by the statistical uncertainties associated with it. Notably, the bins corresponding to higher mass values exhibit significant uncertainties due to sparse event counts. To provide a clearer understanding of the distribution of events within these bins, Fig. 5.27 presents a detailed plot of the event count for 2018.

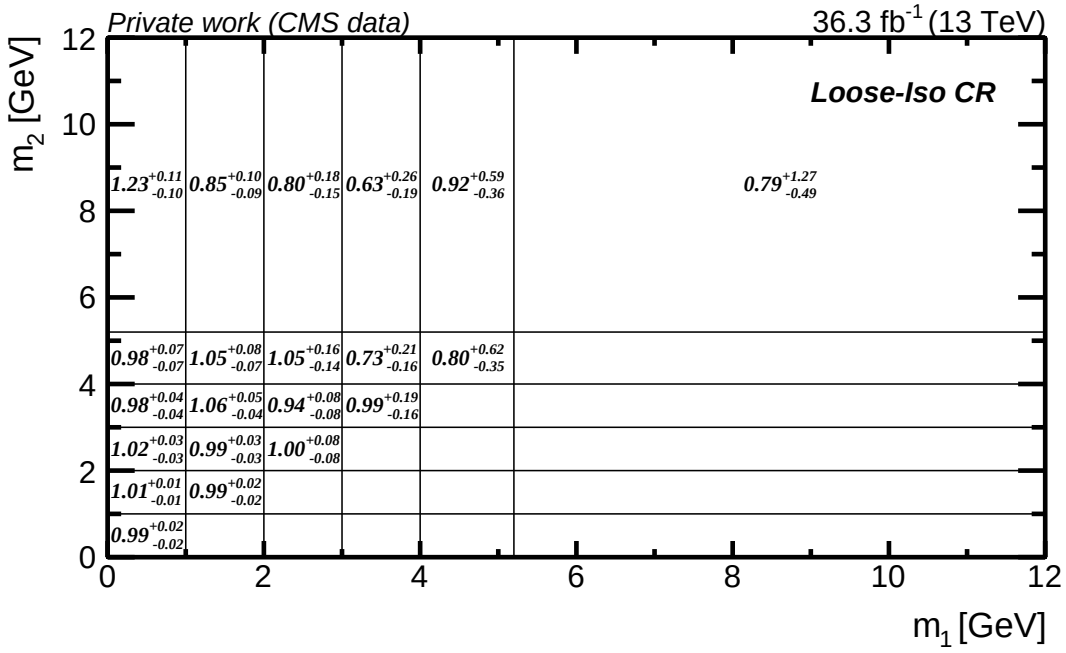

Figure 5.24: The correlation factors $C(i,j)^{\mathrm{CR}}_{\mathrm{data}}$ with statistical uncertainties, estimated with 2016 data.

To estimate $C(i,j)$ within the signal region, it is necessary to adjust the correlation factors derived from the Loose-Iso CR data to account for discrepancies between the control and signal regions. This correction is estimated using simulated background events; directly selecting

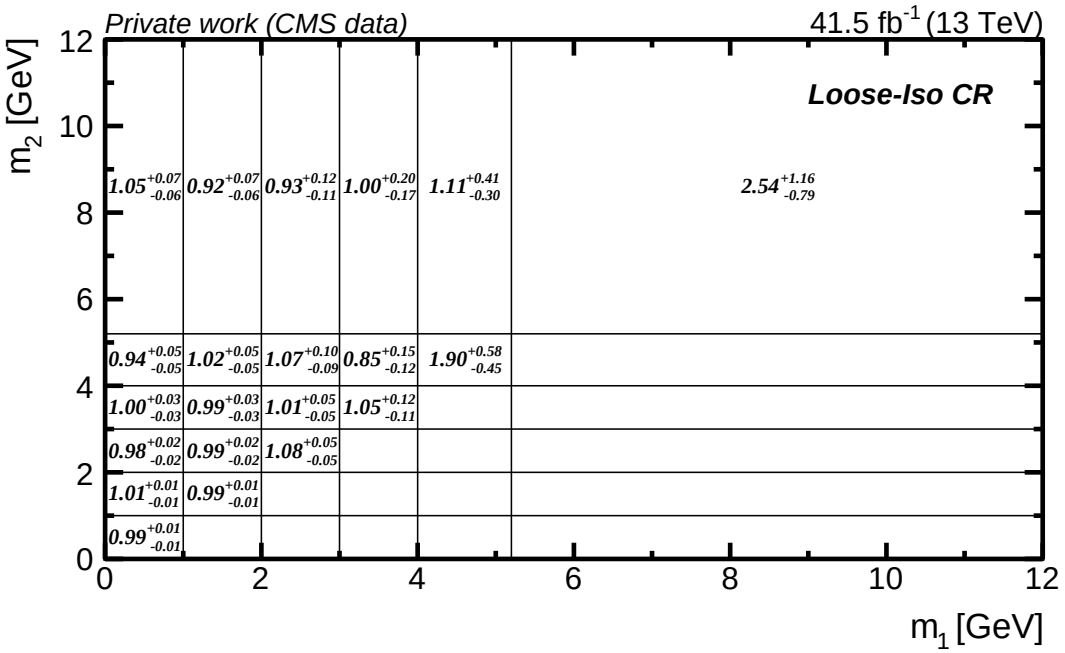

Figure 5.25: The correlation factors $C(i,j)^{\mathrm{CR}}_{\mathrm{data}}$ with statistical uncertainties, estimated with 2017 data.

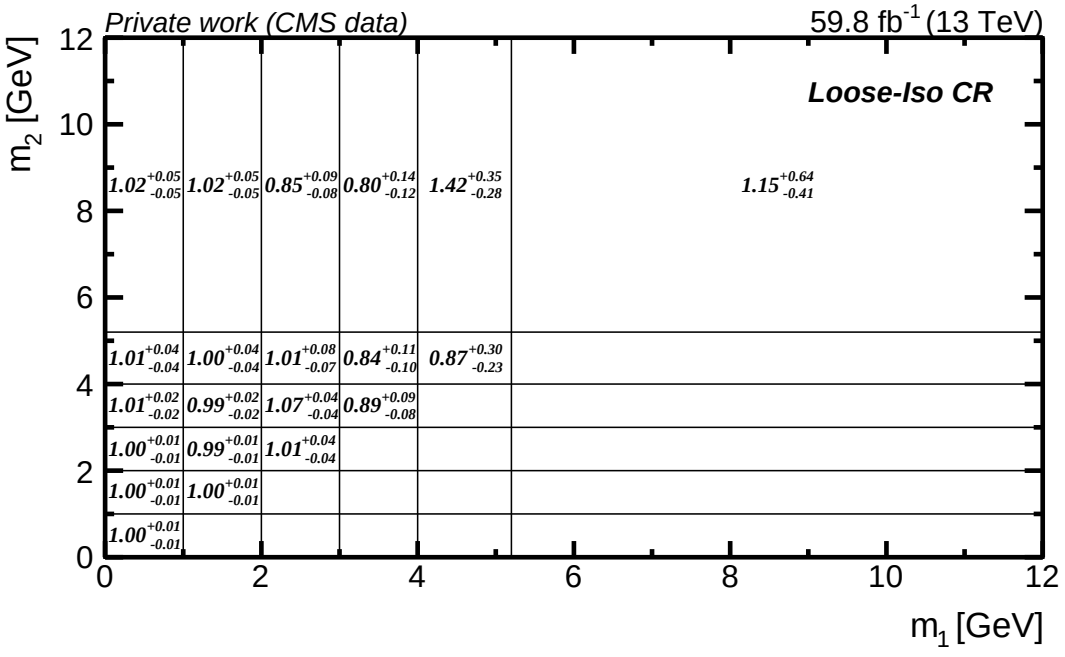

Figure 5.26: The correlation factors $C(i,j)^{\mathrm{CR}}_{\mathrm{data}}$ with statistical uncertainties, estimated with 2018 data.

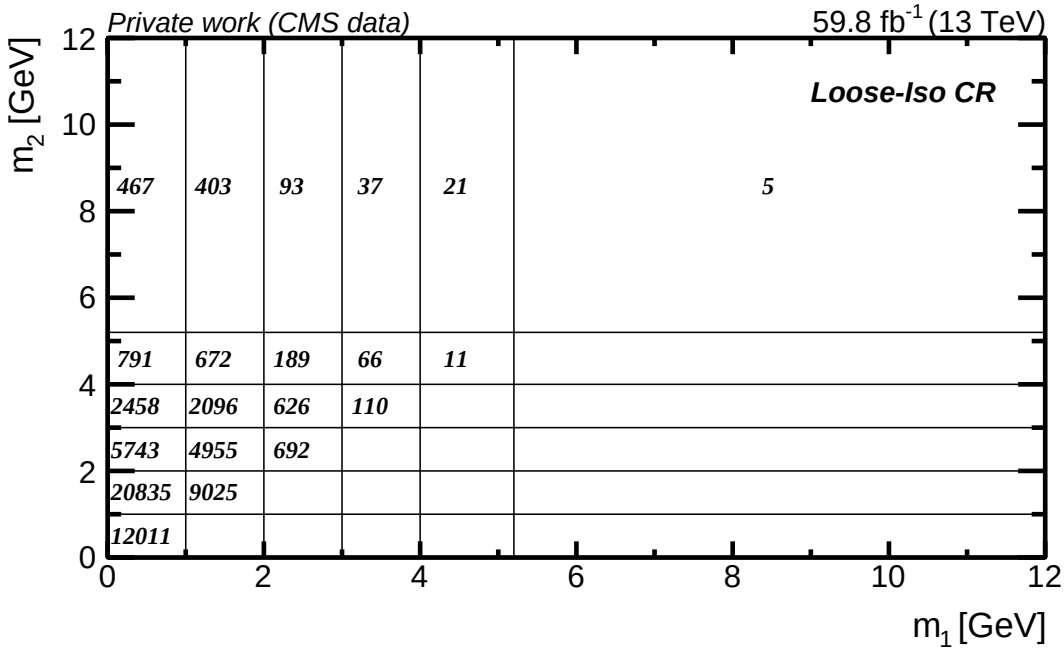

Figure 5.27: Distribution of events across the different bins in the Loose-Iso control region for 2018.

events presents challenges due to the limited statistics in both regions and the significant statistical uncertainties, especially for the simulated QCD multijet events. Consequently, a dedicated Monte Carlo study utilizing QCD samples has been conducted.

This QCD MC study posits that correlations in the invariant masses of muon-track pairs in QCD events stem from the correlations between the flavors of partons initiating the jets that meet the selection criteria in either the SR or the Loose-Iso CR. The study calculates probabilities that a jet containing a muon complies with the criteria for the SR or Loose-Iso CR, considering factors such as the parton flavor, the product of the charges of the muon and parton, the parton momentum, and the ratio of muon-to-parton momentum. Event weights derived from these probabilities are used to construct the normalized distributions $f_{1\mathrm{D}}(i)$ and $f_{2\mathrm{D}}(i,j)$.

For other background processes, which contribute negligibly, no specialized treatment is applied; selections are made directly from the MC to populate the $f_{1\mathrm{D}}(i)$ and $f_{2\mathrm{D}}(i,j)$. The final 1D and 2D distributions, incorporating both weighted QCD and non-QCD events, are then used to compute the correlation factors for both the SR and the Loose-Iso CR.

The test shows good agreement between the data and MC in the Loose-Iso CR, where the signal contribution is expected to be negligible, thereby confirming the reliability of the background model and addressing the challenge posed by limited events.

The correlation factors thus estimated from simulation in the signal region, $C(i,j)^{\mathrm{SR}}_{\mathrm{MC}}$, and the Loose-Iso CR, $C(i,j)^{\mathrm{CR}}_{\mathrm{MC}}$, for the years 2016, 2017 and 2018 are shown in Figures. 5.28, 5.29 and 5.30, respectively.

The correlation factors for the data in the signal region, $C(i,j)^{\mathrm{SR}}$data, are then computed as follows:

$$C(i,j)^{\mathrm{SR}}_{\mathrm{data}} = C(i,j)^{\mathrm{CR}}_{\mathrm{data}} \frac{C(i,j)^{\mathrm{SR}}_{\mathrm{MC}}}{C(i,j)^{\mathrm{CR}}_{\mathrm{MC}}}. \tag{5.7}$$

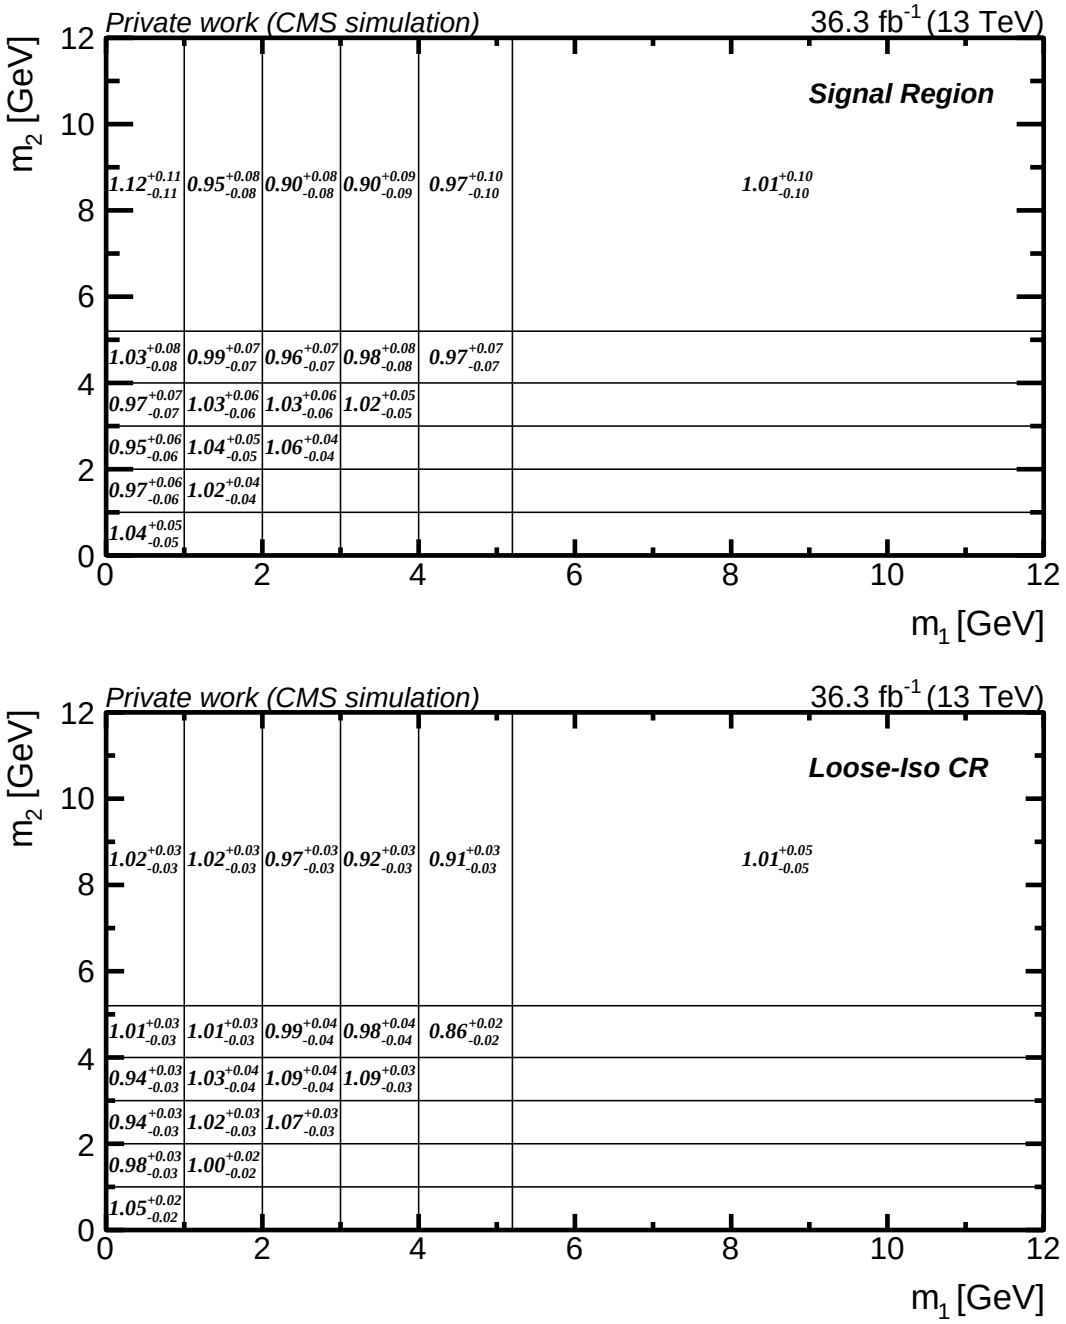

Figure 5.28: The correlation factors $C(i,j)^{\mathrm{SR}}_{\mathrm{MC}}$ (upper) and $C(i,j)^{\mathrm{CR}}_{\mathrm{MC}}$ (lower) with statistical uncertainties for 2016.

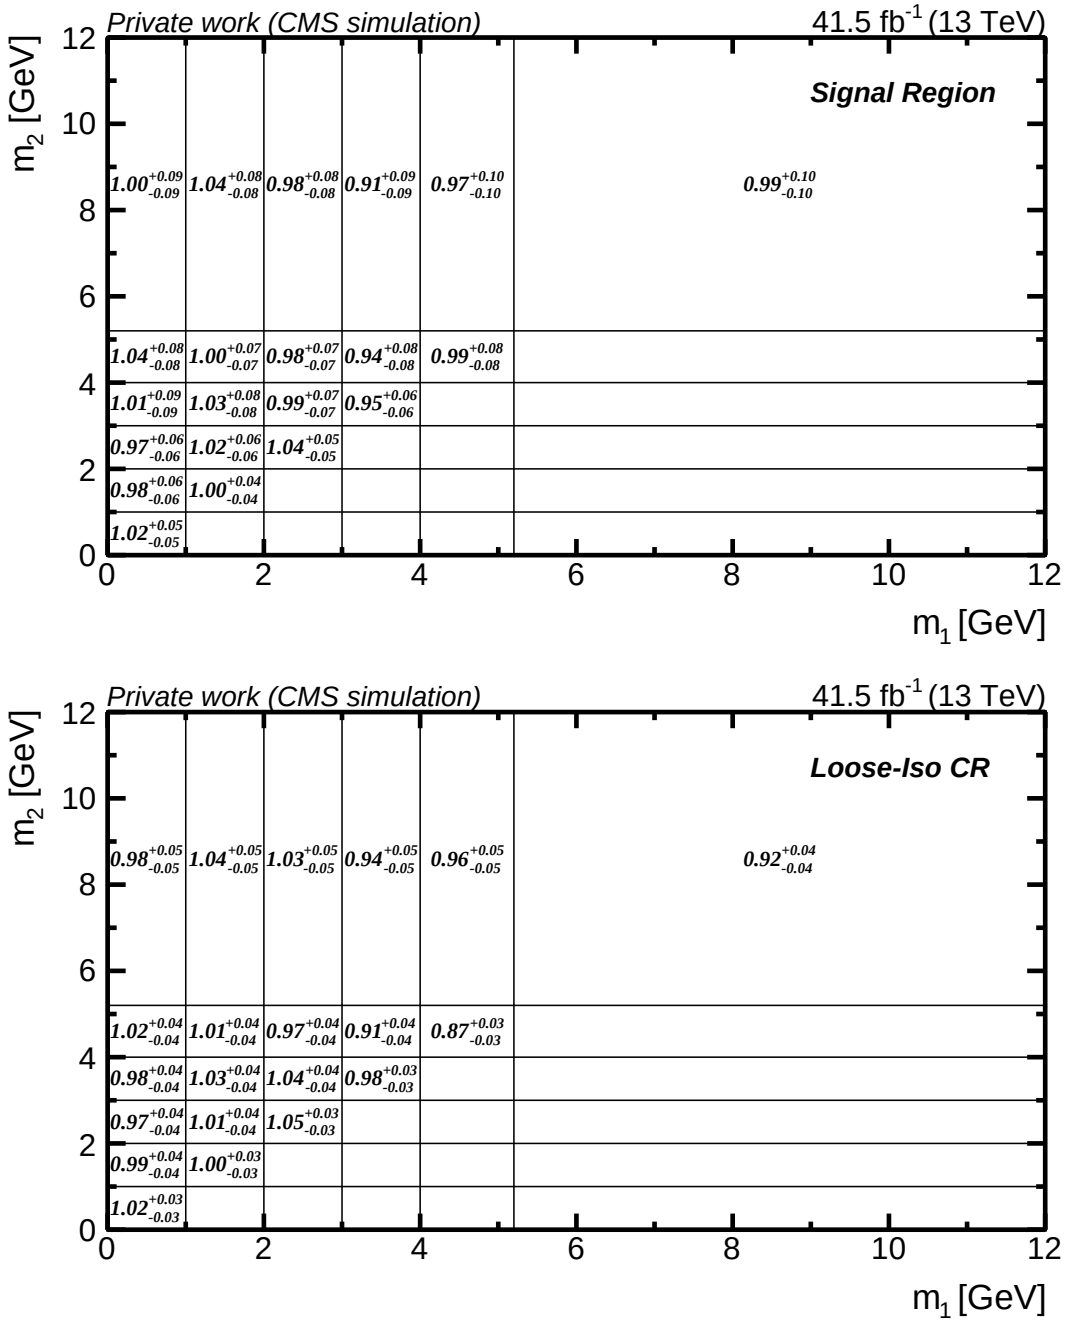

Figure 5.29: The correlation factors $C(i,j)^{\mathrm{SR}}_{\mathrm{MC}}$ (upper) and $C(i,j)^{\mathrm{CR}}_{\mathrm{MC}}$ (lower) with statistical uncertainties for 2017.

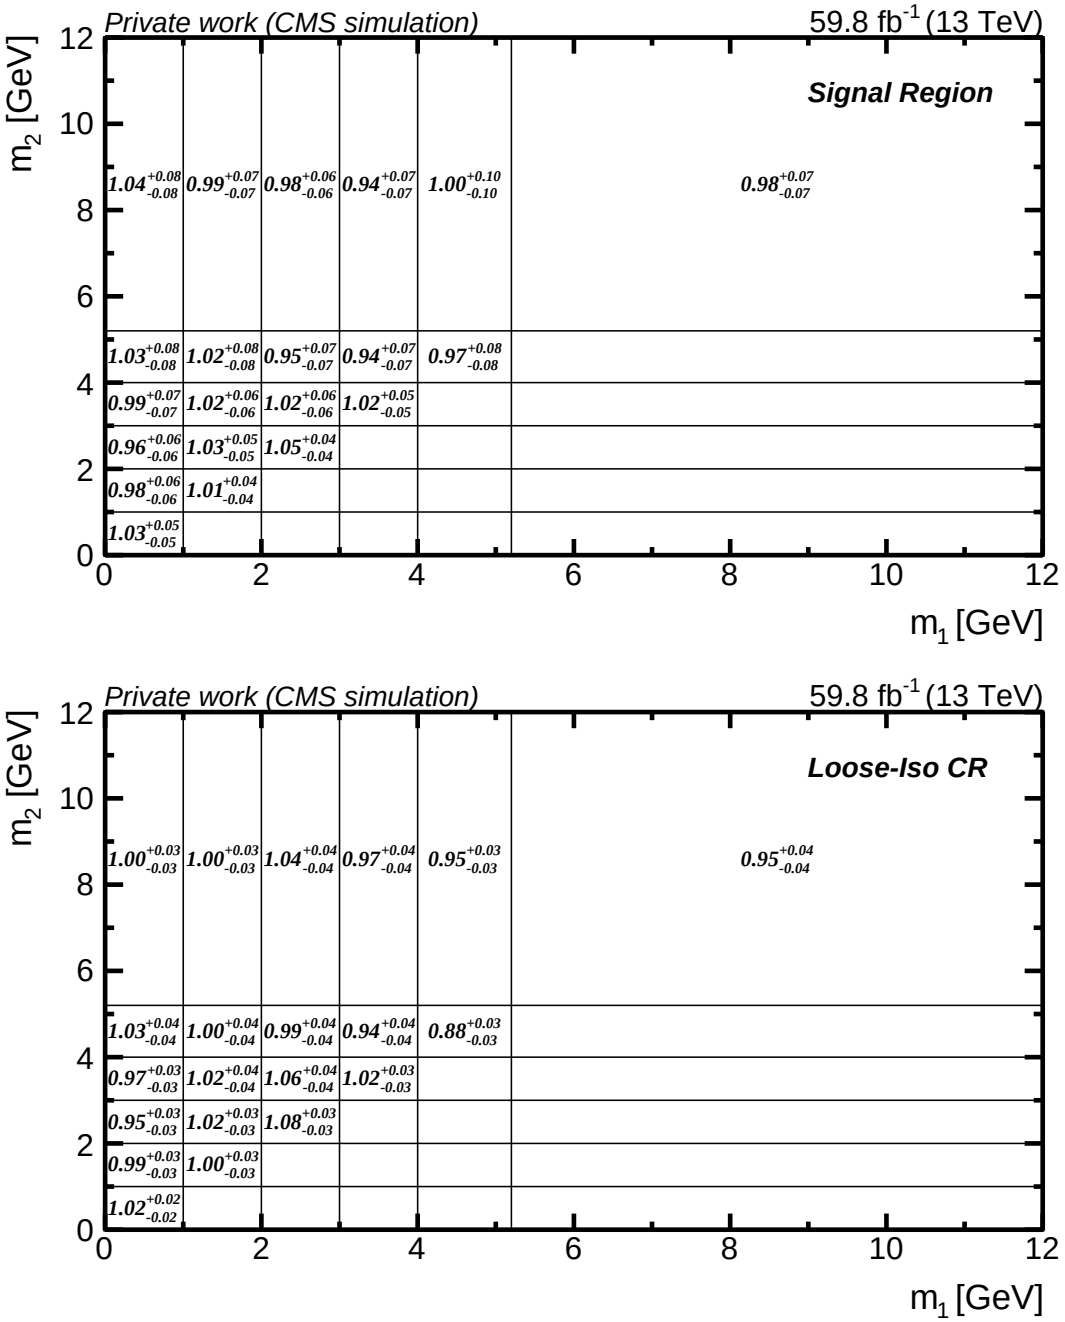

Figure 5.30: The correlation factors $C(i,j)^{\mathrm{SR}}_{\mathrm{MC}}$ (upper) and $C(i,j)^{\mathrm{CR}}_{\mathrm{MC}}$ (lower) with statistical uncertainties for 2018.

Systematic Uncertainties in $C(i, j)$ Estimation

The estimation of correlation factors $C(i, j)$ from simulated events is subject to various systematic uncertainties that can significantly affect the results. These include:

- Uncertainty in parton shower scale: Variations in the modeling of initial and final state radiation (ISR and FSR) during parton showering can influence the reconstructed muon-track masses and their correlations in events where an a_1 candidate is mimicked by a jet. This, in turn, will lead to potential deviations in $C(i, j)$. To evaluate the impact of this uncertainty, the parton shower scale is varied up and down by a factor of 2 and 0.5, respectively.
- Uncertainty in the relative fraction of non-QCD events: The precise contribution of non-QCD events is not well understood, which may introduce inaccuracies in the estimation of $C(i, j)$. To account for this, the yield of non-QCD background is conservatively varied by ±50%. This variation allows the non-QCD fraction to fluctuate between 0% and 4.8%, depending on its initial estimated value.

The extrapolation of the correlation factors from the control region data, utilizing the ratios from Monte Carlo simulations across the signal and control regions, allows for the systematic uncertainties identified for simulations to be effectively transferred to the experimental measurements of the correlation factors in the signal region.

5.10 Systematic Uncertainties

Various sources of systematic uncertainty impact the modeling of the signal and background processes. These uncertainties arise from limited knowledge of the processes, discrepancies between simulations and experimental data, and imperfect understanding of the detector response. Depending on the effect the systematic uncertainties have on the templates, they are categorized as either *normalization* or *shape* uncertainties. Normalization uncertainties affect the total event yield, whereas shape uncertainties also modify the differential distribution of events. Accurately quantifying and incorporating these uncertainties ensures the reliability and precision of the results. The systematic uncertainties considered in this analysis are summarized in Table 5.9 and are described in detail in the subsequent subsections.

Bin-by-bin statistical uncertainties affect both the signal and background templates and arise from the limited size of the data samples in the control regions used for background modeling as well as the simulated samples of background and signal processes. These uncertainties are categorized as a specific type of shape uncertainty and estimated using the Barlow-Beeston-lite method [206], which is detailed in the following chapter.

5.10.1 Uncertainties affecting background

The 2D distribution for the background is influenced by the shape uncertainty arising from the 1D template $f_{1D}(i)$, discussed in Sec. 5.9.1, which accounts for any potential bias introduced by estimating the 1D distribution in the N_{23} CR. The numerical value for these uncertainties varies from bin to bin, affecting the yield between 1-15%. The background shape is further

Uncertainty source	Affected template	Type	Correlation
$f_{1D}(i)$ estimation	Background	shape	No
$C(i,j)$ - ISR	Background	shape	Yes
$C(i,j)$ - FSR	Background	shape	Yes
$C(i,j)$ - non-QCD	Background	shape	Yes
Luminosity	Signal	norm	Yes
Muon ID and trigger efficiency	Signal	norm.	Yes
Track ID/isolation efficiency	Signal	shape	No
Prefiring weights	Signal	norm.	No
Jet Energy Scale	Signal	norm.	No
b tagging efficiency	Signal	norm.	Yes
Bin-by-bin statistical uncertainty	All	shape	No
Theoretical Uncertainties			
μ_R and μ_F variations (acceptance)	Signal	norm.	Yes
μ_R and μ_F variations (cross section)	Signal	norm.	Yes
PDF (acceptance)	Signal	norm.	Yes

Table 5.9: Summary of the sources of uncertainty affecting the signal and background templates in the analysis, categorized into normalization and shape uncertainties. Also mentioned is whether each source is correlated across the data corresponding to different years.

impacted by uncertainties related to the extrapolation of the correlation factors $C(i,j)$ from the Loose-Iso CR to the SR, which stem from systematic variations in the modeling of ISR and FSR and the non-QCD background fraction in the MC simulations as described in Sec. 5.9.2. The ISR, FSR, and non-QCD contributions lead to variations in background yield for individual bins by up to 1%, 2%, and 8%, respectively.

5.10.2 Uncertainties affecting signal distributions

Various sources of uncertainties, both uncorrelated and correlated, can originate from integrated luminosity measurements. Normalization uncertainties of 1.2%, 2.3%, and 2.5% are applied individually for 2016, 2017, and 2018 respectively. Additionally, a correlated uncertainty of 0.27% affects 2017 and 2018, while another 1.3% uncertainty is correlated across all three years.

Muon identification and trigger efficiencies introduce a per-muon uncertainty of 1.5%. Given that two muons are involved in the final state, this translates into a 3% systematic uncertainty in signal acceptance.

Uncertainties from track selection and isolation efficiency, which range from 5-12% per track depending on its p_T, affect the shape of the signal templates and alter the overall signal yield by 10–24%.

Prefiring corrections bring uncertainties ranging from 0.1% to 2.8%, varying by mass point and signal sample. Similarly, jet energy scale measurements introduce up to a 1% uncertainty in the signal yield.

Uncertainties in measuring the b tagging efficiency, as detailed in Ref. [173], are differentiated for heavy-flavor and light-flavor jets in simulated samples. These uncertainties are

categorized into those specific to each data-taking period and those correlated across periods. The uncertainties have the most significant impact on the $t\bar{t}H$ sample, with yield variations ranging from 4 to 6%. For the ggF, VBF, and VH signal samples, the b tagging uncertainties lead to variations in the yield between 0.2–0.5%.

Theoretical uncertainties from missing higher-order corrections in the gluon-gluon fusion process are evaluated using the HQT program. By adjusting the renormalization (μ_R) and factorization (μ_F) scales, the resulting H p_T-dependent k-factors are recomputed and applied to the simulated signal samples. The resulting effect on the signal acceptance is estimated to be approximately 1%. Similar evaluations for the VBF, VH, and $t\bar{t}H$ show uncertainties in acceptance ranging from 1 to 3%, depending on the process and the mass of a_1.

Variations in these scales also affect the cross section calculations for different Higgs boson production mechanisms, introducing uncertainties up to 10%, varying by mass point and signal sample.

Additionally, uncertainties arising from the choice of PDFs are assessed using the HQT program. The nominal k-factors for the H p_T spectrum, calculated with the NNPDF3.1 PDF set [207], show that internal variations within this PDF set modify the signal acceptance by about 1%. Switching to the CTEQ6L1 PDF set [208] alters acceptance by roughly 0.7%. The assigned uncertainty of 1% thus covers these variations. The impact of PDF uncertainties on the acceptance for the VBF, VH, and $t\bar{t}H$ processes is estimated in a similar way, resulting in a 2% uncertainty.

Chapter 6

Results and Interpretation

Contents

6.1 Statistical Methods 125
6.1.1 The Likelihood Function 126
6.1.2 The Maximum Likelihood Method 127
6.1.3 Hypothesis Testing 128
6.2 Analysis Results 130
6.2.1 Maximum-likelihood Fit to the 2D (m_1, m_2) Distribution 131
6.2.2 Goodness-of-fit Test 136
6.2.3 Impacts and Pulls of Nuisance Parameters 138
6.2.4 Model-independent Upper Limits 138
6.3 Interpretation in the context of 2HDM+S Scenarios 144

Following the in-depth description of the analysis setup and techniques in the previous chapter, this chapter is dedicated to presenting the results of the search for light pseudoscalars in exotic decays of the Higgs boson in the 4τ final state. Before presenting the results, the chapter first outlines the statistical methods used to arrive at these findings in Sec 6.1, to provide a clear understanding of how they were derived. Section 6.2 then proceeds with discussing the results, obtained individually for each analyzed year—2016, 2017 and 2018—and also for the Run 2 combination. Finally, the results are interpreted within the context of Two-Higgs-Doublet Models plus a singlet (2HDM+S) scenarios, where the motivation for this search is particularly strong.

6.1 Statistical Methods

In particle physics experiments, the data are subject to stochastic variations due to the quantum mechanical nature of particle collisions and variability in detector responses. This variability necessitates the use of statistical methods to extract meaningful results from the complex data. Within the context of Beyond the Standard Model (BSM) searches, these

methods are utilized not only to probe for potential new signals but also to set upper limits (UL) on the cross-section times branching ratio where no detectable signals are present.

Consequently, this section provides a theoretical overview of the main statistical techniques essential for analyzing and interpreting the experimental results, which will be presented later in this chapter.

6.1.1 The Likelihood Function

The cornerstone of the statistical framework used in all analyses at the Large Hadron Collider (LHC) is the formulation of the likelihood function. This function is defined as the probability density function (pdf) of the observed data, parameterized by a set of model parameters.

As explained in Chapter 5, the analysis detailed in this thesis utilizes binned histograms, where the content of each bin is independent of the others. Each bin can be treated as a distinct counting experiment with an average event rate denoted by r. The probability of observing n events in such an experiment follows the Poisson distribution:

$$P(n|r) = \frac{e^{-r} r^n}{n!} \tag{6.1}$$

However, r is typically unknown and is instead estimated by the expected number of events a over the duration of the experiment. In typical LHC analyses, both signal and background processes contribute to the observed data. Thus, the expected number of events in each bin, a_i is modeled as the sum of expected signal and background events:

$$a_i = \mu s_i + b_i \tag{6.2}$$

Here, s_i represents the number of signal events predicted by the model in bin i, b_i is the corresponding prediction for background events, and μ is the signal strength modifier. The signal strength modifier, usually defined as $\sigma \cdot \mathcal{B}/\sigma_{SM}$, scales the number of expected signal events, providing a measure of how the observed number of signal events compares to the theoretical prediction. A value of $\mu = 0$ corresponds to the scenario of no signal being present.

The likelihood function for a histogram with N bins, where each bin i has n_i observed and a_i expected events, is then defined as the product of the Poisson probabilities for each bin:

$$\mathcal{L}(\text{data}|\mu) = \prod_{i=1}^{N} \frac{e^{-a_i} a_i^n}{n_i!} = \prod_{i=1}^{N} \frac{e^{-(\mu s_i + b_i)} (\mu s_i + b_i)^{n_i}}{n_i!} \tag{6.3}$$

To account for systematic uncertainties that influence the model and, therefore, the expected outcome of the measurement, a set of parameters $\tilde{\theta}$, often called "nuisance" parameters, are introduced. Each source of systematic uncertainty is represented by a different nuisance θ_j, whose nominal value $\tilde{\theta}_j$ is usually determined by dedicated measurements, and its uncertainty is described by the pdf $p(\tilde{\theta}_j|\theta_j)$. The systematic uncertainties change the expected number of events a_i, making it a function of the nuisance parameters $\tilde{\theta} = \{\theta_1, \ldots, \theta_L\}$. Consequently, the likelihood function is modified to integrate these effects:

$$\mathcal{L}(\text{data}|\mu,\tilde{\theta}) = \prod_{i=1}^{N} \frac{e^{-(\mu s_i(\tilde{\theta})+b_i(\tilde{\theta}))}(\mu s_i(\tilde{\theta})+b_i(\tilde{\theta}))^{n_i}}{n_i!} \prod_{j=1}^{L} p(\tilde{\theta}_j|\theta_j) \tag{6.4}$$

The choice of the probability density function for modeling nuisance parameters depends significantly on the nature of the uncertainty. Systematic uncertainties that only affect event yields are commonly modeled using a log-normal distribution:

$$p(\theta) = \frac{1}{\sqrt{2\pi}\ln(\kappa)} \exp\left(-\frac{(\ln(\theta/\tilde{\theta}))^2}{2(\ln(\kappa))^2}\right)\frac{1}{\tilde{\theta}} \tag{6.5}$$

Here κ is the width of the distribution and is chosen to be equal to the measured uncertainty. The log-normal distribution ensures that the estimated event yields are always positive, avoiding unphysical contributions to the expected number of events.

Shape uncertainties, on the other hand, are modeled using a vertical template morphing technique [209]. For each nuisance parameter that affects the shape of the distribution, two additional templates, corresponding to ± 1 standard deviation (σ) variations, are generated. Then, a parameter θ with a Gaussian constraint is added to the likelihood model to smoothly interpolate between the nominal and $\pm 1\sigma$ templates. The Gaussian constraint for the nuisance parameter θ is given by:

$$p(\tilde{\theta}|\theta) = \frac{1}{\sqrt{2\pi}\tilde{\sigma}} \exp\left(-\frac{(\tilde{\theta}-\theta)^2}{2\tilde{\sigma}^2}\right) \tag{6.6}$$

where $\tilde{\sigma}$ is the uncertainty, $\tilde{\theta}$ is the observed value, and θ is the true value of the parameter.

Bin-by-bin uncertainties, arising from limited simulated events in each bin, are a special case of shape uncertainties. These uncertainties are handled using the Barlow-Beeston-lite method [206], which introduces one nuisance parameter per bin, allowing the contents of each bin of each process to vary within its statistical uncertainty.

Lastly, unconstrained or "freely-floating" nuisance parameters are often assigned a flat prior within a realistic range.

6.1.2 The Maximum Likelihood Method

As indicated in the previous section, the likelihood function in high-energy physics analyses depends on the parameter of interest, such as the signal strength modifier μ, and nuisance parameters θ, which account for systematic uncertainties. These model parameters are estimated using a statistical method called maximum likelihood fit.

This method finds the best-fit parameters, $\hat{\mu}$ and $\hat{\theta}$, by maximizing the likelihood function for the observed data or, equivalently, by minimizing the negative log-likelihood (NLL). Minimizing the NLL is computationally more convenient, as it transforms the product of probabilities into a sum. The fit simultaneously adjusts both μ and θ to maximize agreement between the model predictions and the observed data.

Confidence intervals on μ are then determined by profiling the likelihood with respect to the nuisance parameters θ, which incorporates their effects into the estimation of μ. This ensures that the resulting confidence intervals account for systematic uncertainties, yielding robust and unbiased results.

6.1.3 Hypothesis Testing

The hypothesis testing procedure forms the basis of quantifying the presence or absence of a signal in particle physics experiments. This process relies on the distinction between two competing hypotheses: the null hypothesis, which suggests no new signal, and the alternative hypothesis, assuming the presence of a new signal. To discriminate between these hypotheses, a quantity derived from sample data, referred to as the *test statistic*, must be defined.

According to the Neyman-Pearson lemma, the most powerful test for distinguishing between two simple hypotheses with fully specified probability density functions is the likelihood ratio [210]. In high-energy physics, this principle is implemented through the Profile Likelihood Ratio (PLR), defined as:

$$\lambda(\mu) = \frac{\mathcal{L}(\text{data}|\mu, \hat{\hat{\theta}})}{\mathcal{L}(\text{data}|\hat{\mu}, \hat{\theta})}, \tag{6.7}$$

where:

- μ is the specific value of the signal strength modifier being tested,
- θ are the nuisance parameters associated with systematic uncertainties,
- $\hat{\hat{\theta}}$ are the values of the nuisance parameters that maximize the likelihood function for a fixed μ, and
- $\hat{\mu}$ and $\hat{\theta}$ represent the maximum likelihood estimates of μ and θ when both are left unconstrained.

The PLR is often transformed into the test statistic, $\tilde{q}_\mu$, defined as:

$$\tilde{q}_\mu = -2 \ln \lambda(\mu). \tag{6.8}$$

Once the choice of the test statistic $\tilde{q}_\mu$ is made, then the probability of observing a test statistic as extreme as, or more extreme than, the one calculated, $\tilde{q}_\mu^{\text{obs}}$ can be determined. This probability, referred to as the p-value, is computed for the signal-plus-background hypothesis (p_μ) as:

$$p_\mu = P(\tilde{q}_\mu \geq \tilde{q}_\mu^{\text{obs}}|s+b) = \int_{\tilde{q}_\mu^{\text{obs}}}^{\infty} f(\tilde{q}_\mu|\mu, \hat{\theta}_\mu^{\text{obs}})\, d\tilde{q}_\mu, \tag{6.9}$$

where:

- $\tilde{q}_\mu^{\text{obs}}$ is the observed value of the test statistic for a hypothesized signal strength μ,

- $\hat{\theta}_{\mu}^{\text{obs}}$ are the best-fit nuisance parameter values for a fixed μ, and
- $f(\tilde{q}_{\mu}|\mu, \hat{\theta}_{\mu}^{\text{obs}})$ is the pdf of $\tilde{q}_{\mu}$ under the signal-plus-background hypothesis and can be obtained using pseudo-data, or toys, generated from the signal-plus-background expectation.

For the background-only hypothesis ($\mu = 0$), a related quantity is often computed to assess the significance of rejecting this hypothesis:

$$1 - p_b = P(\tilde{q}_{\mu} \geq \tilde{q}_{\mu}^{\text{obs}}|b) = \int_{\tilde{q}_{\mu}^{\text{obs}}}^{\infty} f(\tilde{q}_{\mu}|0, \hat{\theta}_0^{\text{obs}})\, d\tilde{q}_{\mu}, \tag{6.10}$$

where:

- $\tilde{q}_{\mu}^{\text{obs}}$ is the observed value of the test statistic for a signal strength $\mu = 0$,
- $\hat{\theta}_0^{\text{obs}}$ are the best-fit nuisance parameter values for $\mu = 0$, and
- $f(\tilde{q}_{\mu}|0, \hat{\theta}_0^{\text{obs}})$ is the pdf of $\tilde{q}_{\mu}$ under the background-only hypothesis and can also be obtained using pseudo-data, or toys, generated from the background-only expectations.

The p-values described above provide a measure of how well the data aligns with each hypothesis. A smaller p_{μ} indicates that the data are less compatible with the assumed signal strength μ. Conversely, a larger p_b (or smaller $1 - p_b$) suggests consistency with the background-only hypothesis.

6.1.3.1 Setting Upper Limits

In searches for BSM processes, setting upper limits on the signal strength becomes essential when no new phenomena are observed. These limits are determined using the *limit test statistic*, defined as follows:

$$\tilde{q}_{\mu} \equiv \begin{cases} -2 \ln \lambda(\mu) & \text{if } \hat{\mu} \leq \mu \\ 0 & \text{if } \hat{\mu} > \mu \end{cases} \tag{6.11}$$

The requirement on μ ensures that upward fluctuations of the data, if larger than expected for a signal strength of μ, are not mistakenly regarded as evidence against the signal hypothesis.

The limit test statistics can then be used to compute the observed upper limits following the modified frequentist approach employing the CL_s criterion [211]. The value of CL_s is defined as:

$$\text{CL}_s(\mu) \equiv \frac{\text{CL}_{s+b}(\mu)}{\text{CL}_b(\mu)} \equiv \frac{p_{\mu}}{1 - p_b} \tag{6.12}$$

In practical terms, a signal strength μ is considered to be excluded at a certain confidence level (CL) α if $\text{CL}_s < 1 - \alpha$. Typically, a CL of 95% is adopted for this purpose, meaning that μ values yielding $\text{CL}_s < 0.05$ are excluded.

Another crucial component of limit setting is the calculation of expected limits. Expected limits are needed to provide a statistical benchmark that indicates where the limits might fall if only background processes were present. They are calculated by simulating a large number of background-only pseudo-experiments, which mimic possible outcomes under the background-only hypothesis. Each simulation involves recalculating the limit test statistic, $\tilde{q}_\mu$, and observing how it varies across simulations. The median of these simulations provides the expected limit, while the $16^{\text{th}}/84^{\text{th}}$ and $2.5^{\text{th}}/97.5^{\text{th}}$ percentiles determine the $\pm1\sigma$ and $\pm2\sigma$ bands, respectively. These bands offer insight into the potential variability in the exclusion limits due to statistical fluctuations.

The method of generating toy Monte Carlo simulations for observed or expected limits is computationally demanding, making it impractical for routine experimental physics use. The asymptotic approximation method [212] offers a practical alternative by leveraging Wilks' theorem [213], which predicts that test statistics will converge to a chi-squared distribution as sample sizes increase. This simplifies the calculation of p-values and eliminates the need for extensive toy data simulations. For setting expected limits, the Asimov dataset is invaluable; in this dataset, each observation in the bins matches exactly the properties predicted by the model. The dataset then enables the precise calculation of test statistics and the derivation of expected limits and confidence intervals.

6.2 Analysis Results

Following a comprehensive theoretical discussion of statistical methods in the preceding section, this section applies these methodologies to the experimental search for a pair of light pseudoscalars in final states with four tau leptons. This analysis was performed using data collected during the LHC Run 2 over the years 2016 to 2018. All necessary inputs for statistical inference—including the signal model, background model, and systematic uncertainties—have already been discussed in the previous chapter. The analysis progresses by performing a maximum likelihood fit to the 2D (m_1, m_2) distribution to extract the signal strength modifier $\mu = \sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1a_1} \to 4\tau)/\sigma_{SM}$ and subsequently set upper limits on μ.

The statistical calculations done in this analysis are performed with the COMBINE TOOL package [214], which was originally developed for the combination of results on the searches for the Higgs boson by the CMS experiment. This tool is powerful and versatile, enabling the performance of fits to extract signal strength, conducting goodness-of-fit tests, facilitating fit diagnostics, and determining the influence of nuisance parameters on the parameters of interest.

To maintain the integrity of the analysis while developing search methodologies, the data in the signal region were initially blinded. Instead, data from control regions and Asimov data sets were employed to refine the analytical approach. This enables rigorous testing and enhancement of methods under controlled, bias-free conditions. Upon completion and validation of these refined methods, the real data were unblinded.

This section presents the unblinded results of the $\mathrm{H} \to \mathrm{a_1a_1}$ search, obtained for each analyzed year individually as well as for the full Run 2 combination.

6.2.1 Maximum-likelihood Fit to the 2D (m_1, m_2) Distribution

The signal is extracted using a binned maximum-likelihood fit applied to the (m_1, m_2) distribution. The normalizations of both signal and background are allowed to float freely in the fit. This would detect any potential excesses that the signal-plus-background model might explain and adjust the inherently unknown background normalization. Systematic uncertainties that impact the normalization of the signal templates are addressed by incorporating nuisance parameters with log-normal distributions into the model. Additionally, shape-altering systematic uncertainties are modeled using nuisance parameters, which continuously morph the signal or background template shapes, and are assigned Gaussian prior probability density functions. To account for bin-by-bin statistical uncertainties, the Barlow-Beeston-lite method is utilized.

For each mass hypothesis of the a_1 boson, the (m_1, m_2) distribution undergoes two distinct fits: one using only the background template and another combining both the signal and background templates. The results from these fits suggest that the background distribution consistently aligns with the observed event counts, whereas the estimated signal contribution remains minimal across all probed masses. Consequently, no significant event excess over the background expectation is detected.

Figures 6.1–6.4 illustrates the (m_1, m_2) distribution for the individual years and for Run 2. For visualization purposes, the background template has been normalized by fitting the data under the background-only hypothesis. Signal expectations for m_{a_1} values of 5, 8, 12, and 15 GeV are also presented, assuming the Standard Model production rate for the H boson and a 5% branching fraction for the $H \to a_1 a_1 \to 4\tau$ decays. It is apparent that there are no significant deviations in the data from the background predictions in most bins, although a slight statistical fluctuation is noted in a few bins, indicating a minor deficit of data.

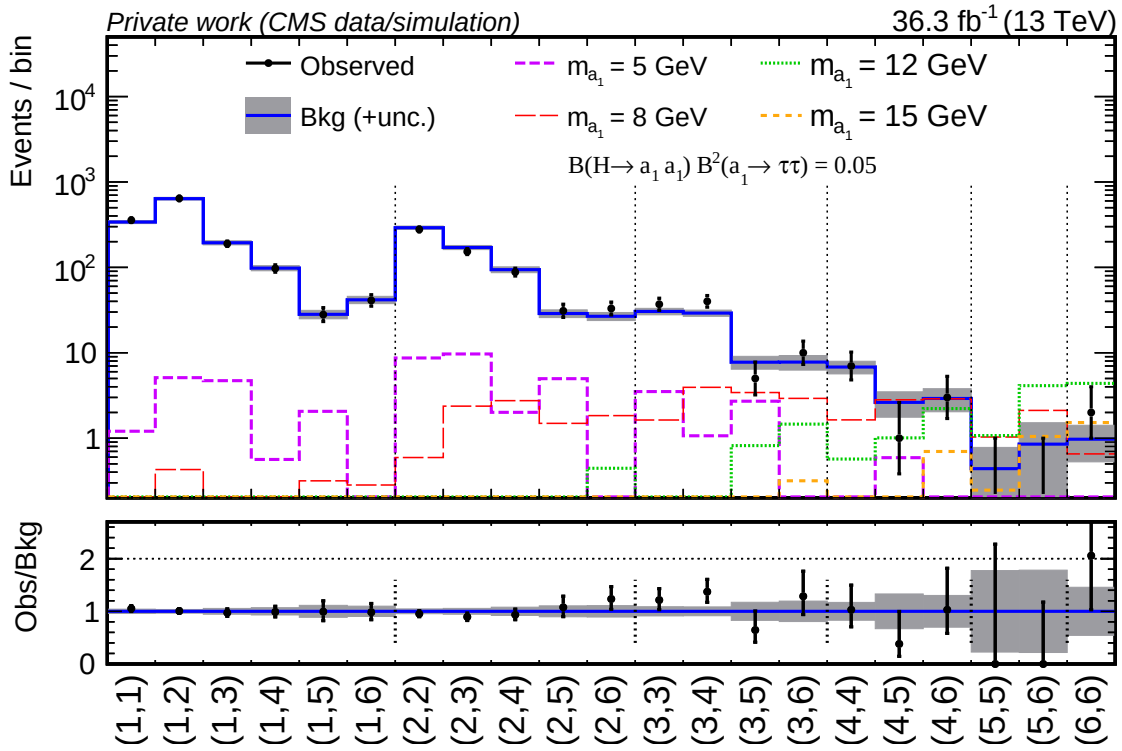

Figure 6.1: The (m_1, m_2) in one-row distribution used to extract the signal for 2016. The observed number of events is represented by data points with error bars. The background with its uncertainty is shown as the blue histogram with the shaded error band. The normalization for the background is obtained by fitting the observed data under the background-only hypothesis. Signal expectations for the 4τ and $2\mu2\tau$ final states are shown as dashed histograms for the mass hypotheses $m_{a_1} = 5$, 8, 12, and 15 GeV. The relative normalization of the 4τ and $2\mu2\tau$ final states are given by Eq. (5.2) as explained in Section 5.8. The signal normalization is computed assuming that the H boson is produced in pp collisions with a rate predicted by the SM and decays into $a_1a_1 \to 4\tau$ final state with the branching fraction of 5%. The lower plot shows the ratio of the observed data events to the expected background yield in each bin of the (m_1, m_2) distribution.

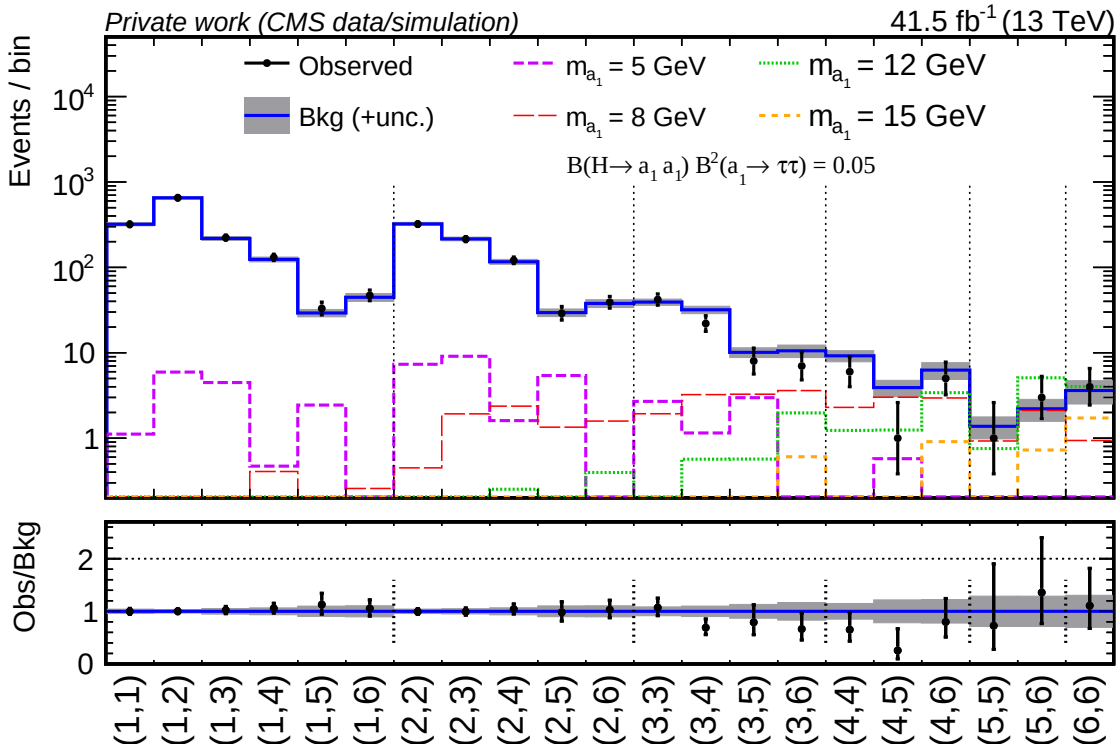

Figure 6.2: The (m_1, m_2) in one-row distribution used to extract the signal for 2017. The observed number of events is represented by data points with error bars. The background with its uncertainty is shown as the blue histogram with the shaded error band. The normalization for the background is obtained by fitting the observed data under the background-only hypothesis. Signal expectations for the 4τ and $2\mu 2\tau$ final states are shown as dashed histograms for the mass hypotheses $m_{a_1} = 5$, 8, 12, and 15 GeV. The relative normalization of the 4τ and $2\mu 2\tau$ final states are given by Eq. (5.2) as explained in Section 5.8. The signal normalization is computed assuming that the H boson is produced in pp collisions with a rate predicted by the SM and decays into $a_1 a_1 \to 4\tau$ final state with the branching fraction of 5%. The lower plot shows the ratio of the observed data events to the expected background yield in each bin of the (m_1, m_2) distribution.

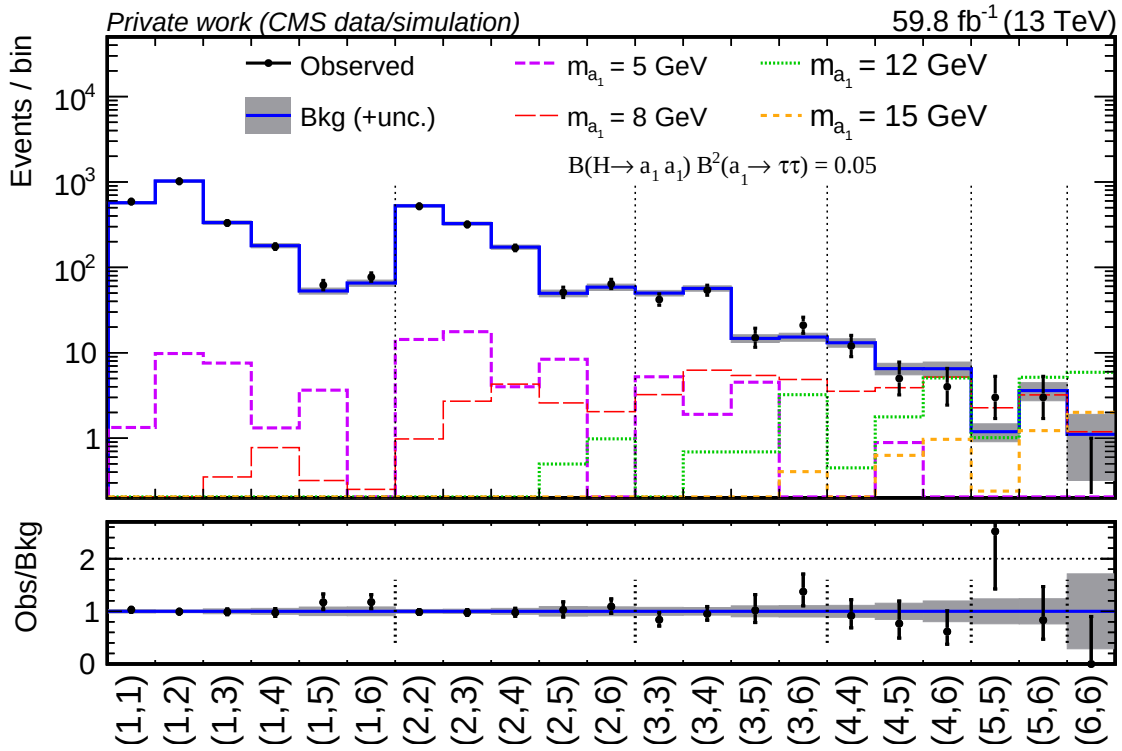

Figure 6.3: The (m_1, m_2) in one-row distribution used to extract the signal for 2018. The observed number of events is represented by data points with error bars. The background with its uncertainty is shown as the blue histogram with the shaded error band. The normalization for the background is obtained by fitting the observed data under the background-only hypothesis. Signal expectations for the 4τ and $2\mu2\tau$ final states are shown as dashed histograms for the mass hypotheses $m_{a_1} = 5$, 8, 12, and 15 GeV. The relative normalization of the 4τ and $2\mu2\tau$ final states are given by Eq. (5.2) as explained in Section 5.8. The signal normalization is computed assuming that the H boson is produced in pp collisions with a rate predicted by the SM and decays into $a_1a_1 \to 4\tau$ final state with the branching fraction of 5%. The lower plot shows the ratio of the observed data events to the expected background yield in each bin of the (m_1, m_2) distribution.

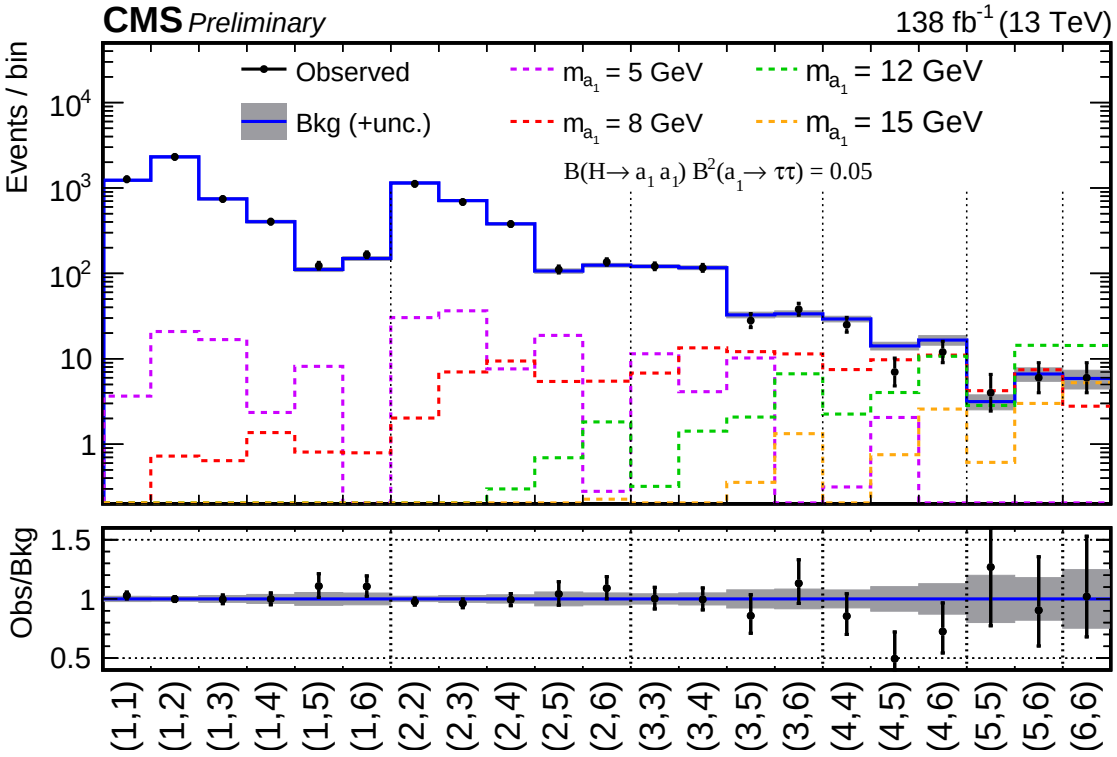

Figure 6.4: The (m_1, m_2) in one-row distribution used to extract the signal for the Run 2 combination. The observed number of events is represented by data points with error bars. The background with its uncertainty is shown as the blue histogram with the shaded error band. The normalization for the background is obtained by fitting the observed data under the background-only hypothesis. Signal expectations for the 4τ and $2\mu2\tau$ final states are shown as dashed histograms for the mass hypotheses m_{a_1}= 5, 8, 12, and 15 GeV. The relative normalization of the 4τ and $2\mu2\tau$ final states are given by Eq. (5.2) as explained in Section 5.8. The signal normalization is computed assuming that the H boson is produced in pp collisions with a rate predicted by the SM and decays into $a_1a_1 \to 4\tau$ final state with the branching fraction of 5%. The lower plot shows the ratio of the observed data events to the expected background yield in each bin of the (m_1, m_2) distribution.

6.2.2 Goodness-of-fit Test

Goodness-of-fit tests are also performed to evaluate how well the statistical models describe observed data. For this analysis, the test is performed using the saturated model as a reference. The saturated model assumes that the observed number of events in each bin exactly matches the predicted number, making it a fully unconstrained benchmark. The goodness-of-fit test compares the negative log-likelihood of the fitted model to that of the saturated model, with the difference serving as the test statistic. The observed test statistic is then compared to its expected distribution under the null hypothesis, which is typically derived using asymptotic approximations or toy Monte Carlo simulations. This comparison yields a p-value, which indicates the probability of observing a discrepancy as large as (or larger than) the one measured, assuming the model is correct. A low p-value suggests poor compatibility between the data and the model.

The results of the goodness-of-fit (GOF) test for the signal-plus-background model, evaluated at a few representative signal mass points, are shown in Fig. 6.5, corresponding to the full Run 2 datasets. In addition, the p-values for these tests, as well as the tests with data from individual years (2016, 2017, and 2018), are summarized in Table 6.5. In all cases, a good compatibility between the model and the observed data is observed. Based on the results of the maximum likelihood fit, GOF tests were performed using the background-only model to test the robustness of this distribution when trying to describe the data. The corresponding results are also presented in Table 6.1. These results indicate that the background-only models are also well-constructed and provide an adequate description of the data.

Period	signal-plus-background				background-only
	5 GeV	8 GeV	12 GeV	15 GeV	-
2016	0.20	0.18	0.17	0.18	0.21
2017	0.75	0.75	0.74	0.74	0.77
2018	0.41	0.41	0.44	0.42	0.45
Run 2	0.45	0.46	0.42	0.45	0.45

Table 6.1: The results of the goodness-of-fit tests for both the signal-plus-background and background-only models, presented in terms of p-values. The p-values were calculated for data collected in 2016, 2017, 2018, and the full Run 2 dataset. For the signal-plus-background hypothesis, the tests were performed at representative signal mass points.

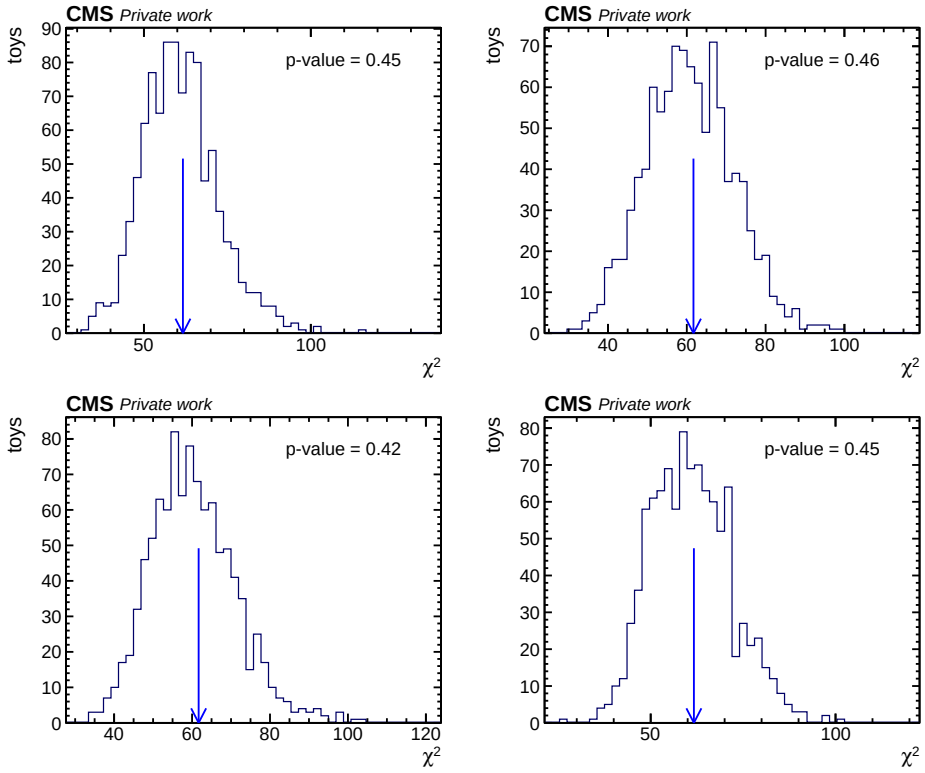

Figure 6.5: Results of the goodness-of-fit tests performed using full Run 2 data for the signal-plus-background model, corresponding to four representative a_1 mass points: 5 GeV (top left), 8 GeV (top right), 10 GeV (bottom left) and 15 GeV (bottom right). The p-values obtained are also shown in each case.

6.2.3 Impacts and Pulls of Nuisance Parameters

Another important check performed is assessing the influence of nuisance parameters on the result of the analysis by evaluating the so-called pulls and impacts. The impact of a nuisance parameter is defined as the shift Δr in the parameter of interest (here, the signal strength) induced by varying a nuisance parameter within its uncertainty, while the other nuisance parameters are fixed to their profiled maximum likelihood estimate. A pull, on the other hand, quantifies how much a nuisance parameter deviates from its nominal (prior) value after the fit to the data. For a given nuisance parameter θ, it is calculated as

$$\text{pull}(\theta) = \frac{\hat{\theta} - \theta_0}{\Delta\theta} \tag{6.13}$$

where $\hat{\theta}$, θ_0, and $\Delta\theta$ are the postfit value, the prefit value, and the prefit uncertainty of the nuisance parameter, respectively. Pulls help identify significant deviations that may suggest tensions between the data and the model or poorly modeled uncertainties, while impacts ensure that the observed effects are not dominated by individual nuisance parameters.

Figure 6.6 shows the pulls and impacts of the 30 most significant nuisance parameters, derived from a maximum likelihood fit to the data for the a_1 mass hypotheses of 5 GeV and 15 GeV, based on the Run 2 combination. The largest contributions to the uncertainty in the signal strength modifier are driven by bin-by-bin statistical uncertainties, labeled as `prop_binchX_binY` in the plots. Other nuisance parameters associated with the background, such as the background normalization (`CMS_haa4t_bkgNorm_201Z`), also exhibit substantial impacts. None of the nuisance parameters were found to be significantly constrained for any mass hypothesis. Most nuisance parameters have pulls within $\pm1\sigma$, with pulls exceeding $\pm2\sigma$ observed only for a few isolated parameters, all of which are related to bin-by-bin statistical uncertainties.

6.2.4 Model-independent Upper Limits

Following the maximum likelihood fits which indicated the absence of any signal, model-independent upper limits at 95% CL on the product of the cross section and branching fraction, $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1)\mathcal{B}^2(\mathrm{a}_1 \to \tau\tau)$, relative to the inclusive SM H boson production cross section, σ_{SM} are set. The limit setting is performed using the modified frequentist $\mathrm{CL_s}$ criterion.

Figures 6.7, 6.8, and 6.9 show the results obtained for 2016, 2017, and 2018, respectively. The observed limits for all mass points are found to be within or near the $\pm2\sigma$ margin of the expected limits.

The limits obtained for the Run 2 combination are shown in Fig. 6.10. The observed limits range from 0.007 at $m_{\mathrm{a}_1} = 11\,\mathrm{GeV}$ to 0.079 at $m_{\mathrm{a}_1} = 4\,\mathrm{GeV}$. The expected upper limits range from 0.011 at $m_{\mathrm{a}_1} = 11\,\mathrm{GeV}$ to 0.066 at $m_{\mathrm{a}_1} = 4\,\mathrm{GeV}$. The results for all mass-points are summarized in Appendix G. The observed limits are compatible with the expected limits within two standard deviations in the entire range of m_{a_1}considered.

Incorporating the $\mathrm{a}_1\mathrm{a}_1 \to \mu\mu\tau\tau$ channel into the signal model consistently improves both the expected and observed limits across all mass points, with particularly notable improvements

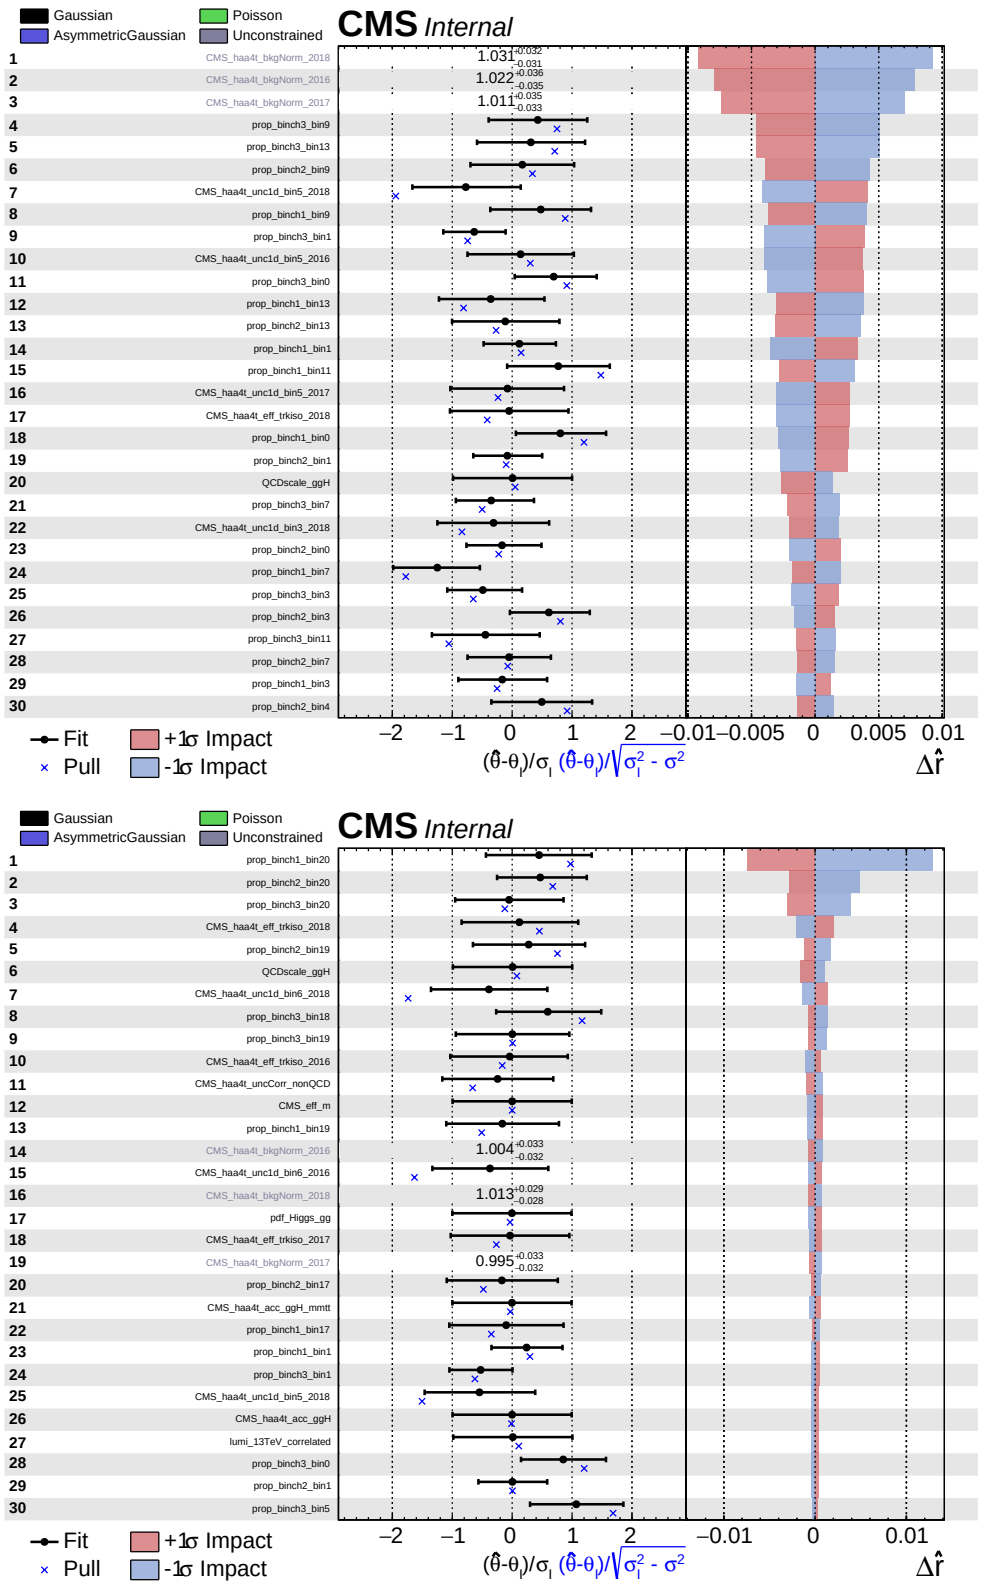

Figure 6.6: The pulls and impacts of the 30 most significant nuisance parameters, for the a_1 mass hypotheses of 5 GeV (top) and 15 GeV (bottom), corresponding to the Run 2 combination.

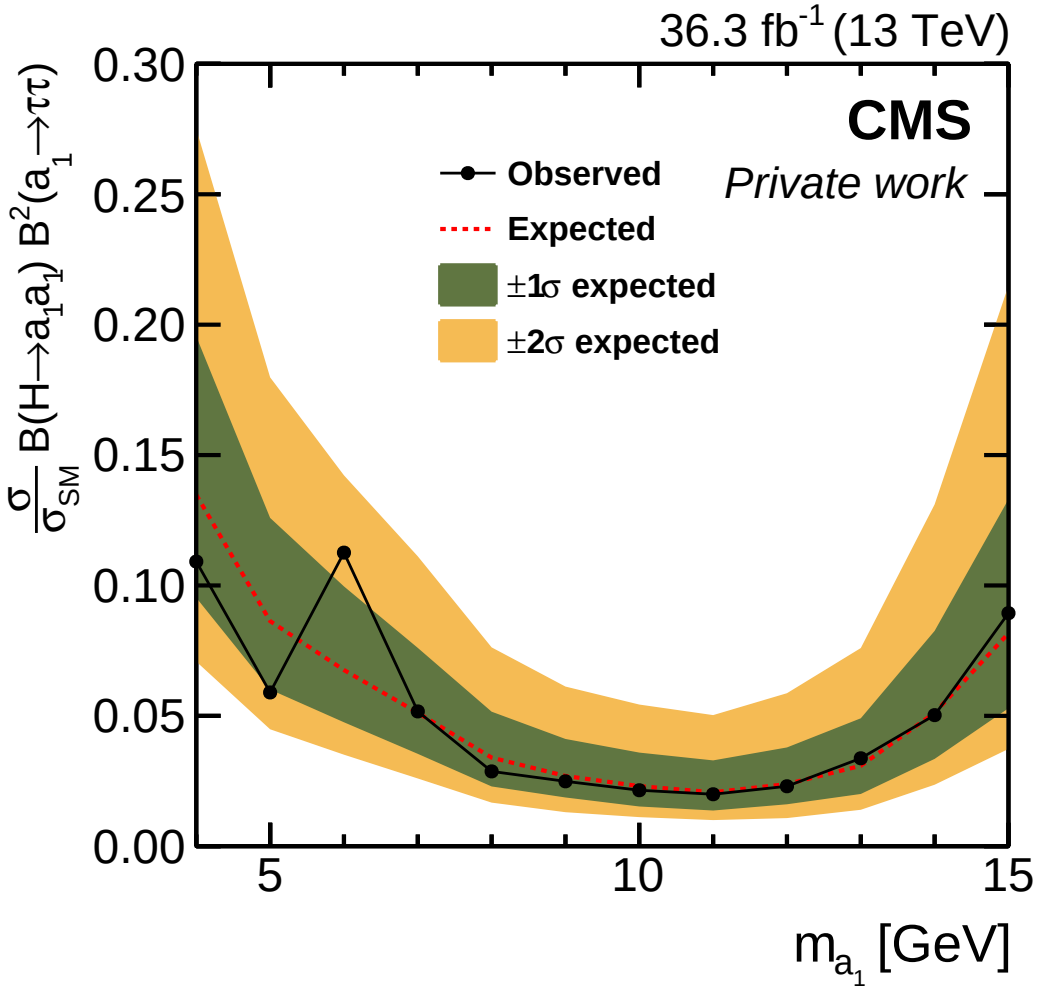

Figure 6.7: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM. These limits are plotted as a function of the $\mathrm{a_1}$ mass for the year 2016. The green and yellow bands represent the 68% and 95% confidence intervals of the limits expected under the background-only hypothesis.

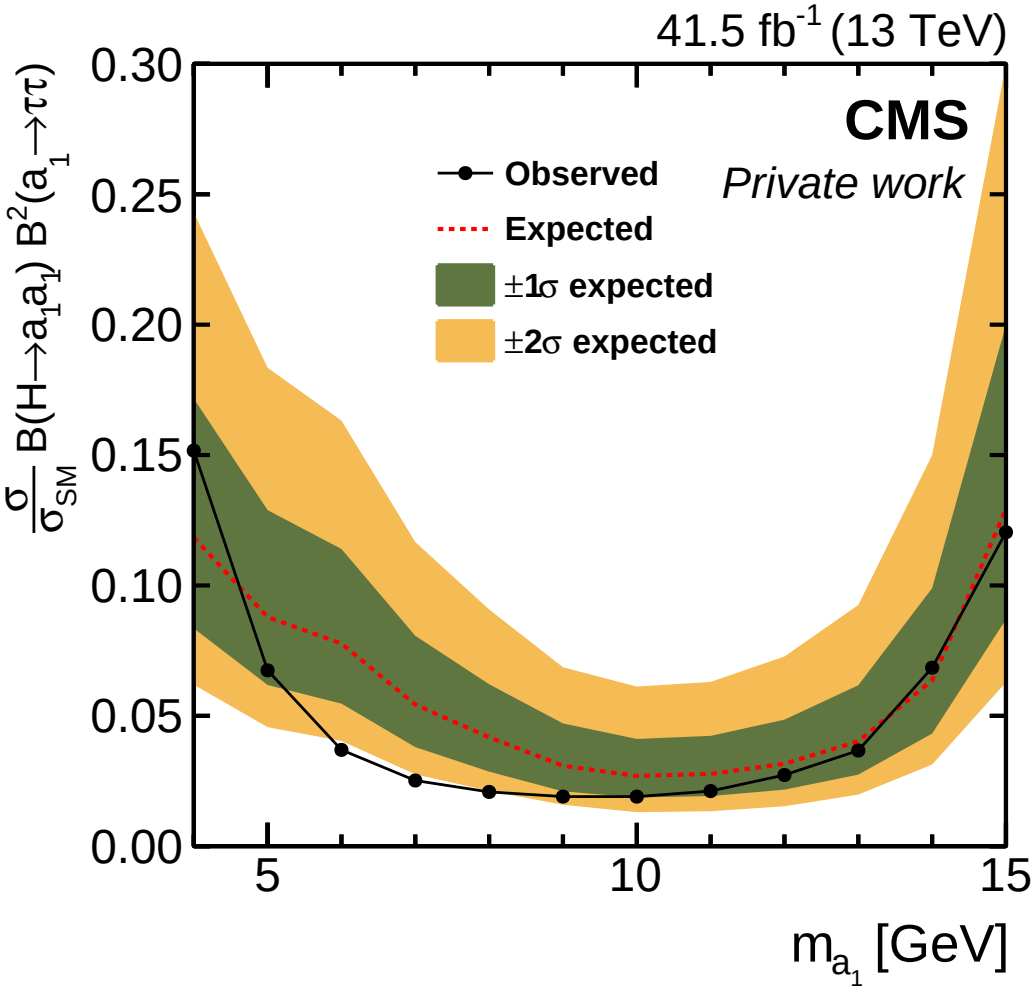

Figure 6.8: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM. These limits are plotted as a function of the $\mathrm{a_1}$ mass for the year 2017. The green and yellow bands represent the 68% and 95% confidence intervals of the limits expected under the background-only hypothesis.

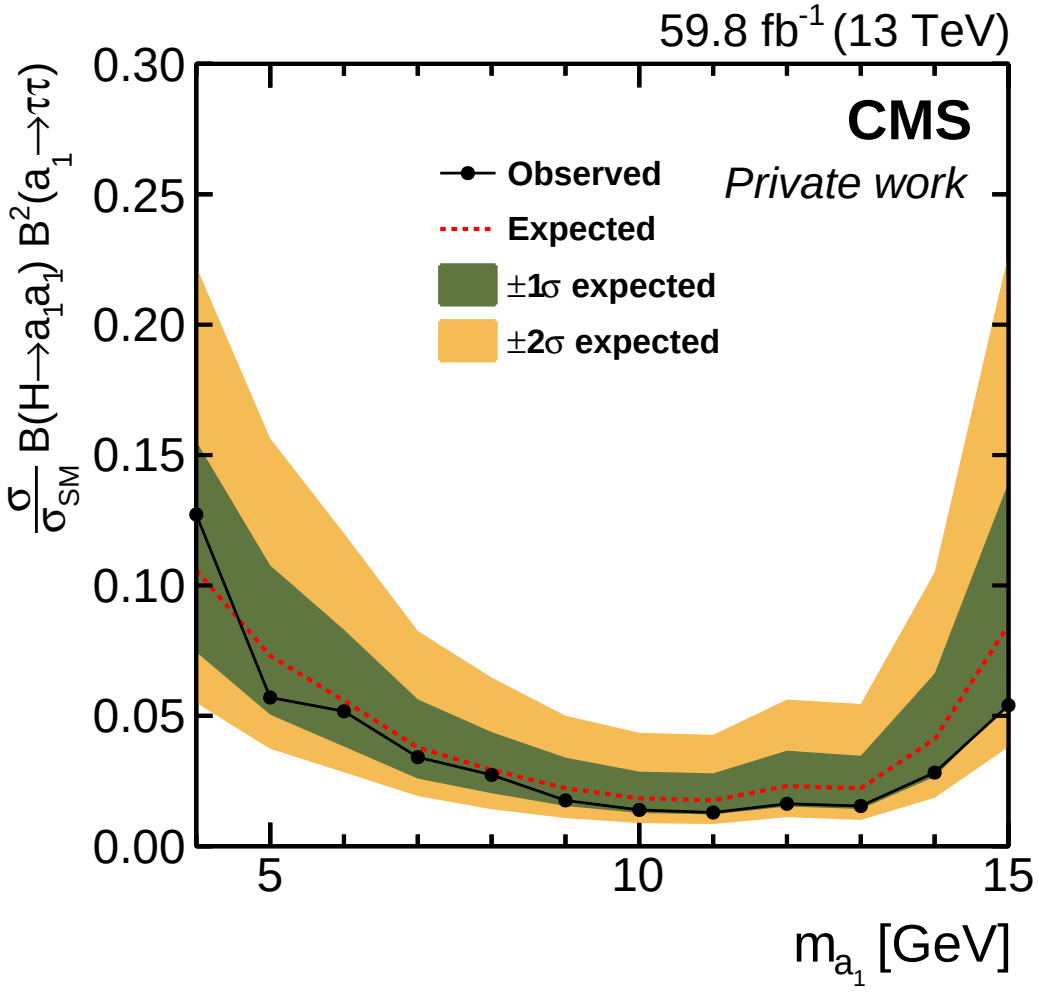

Figure 6.9: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM. These limits are plotted as a function of the $\mathrm{a_1}$ mass for the year 2018. The green and yellow bands represent the 68% and 95% confidence intervals of the limits expected under the background-only hypothesis.

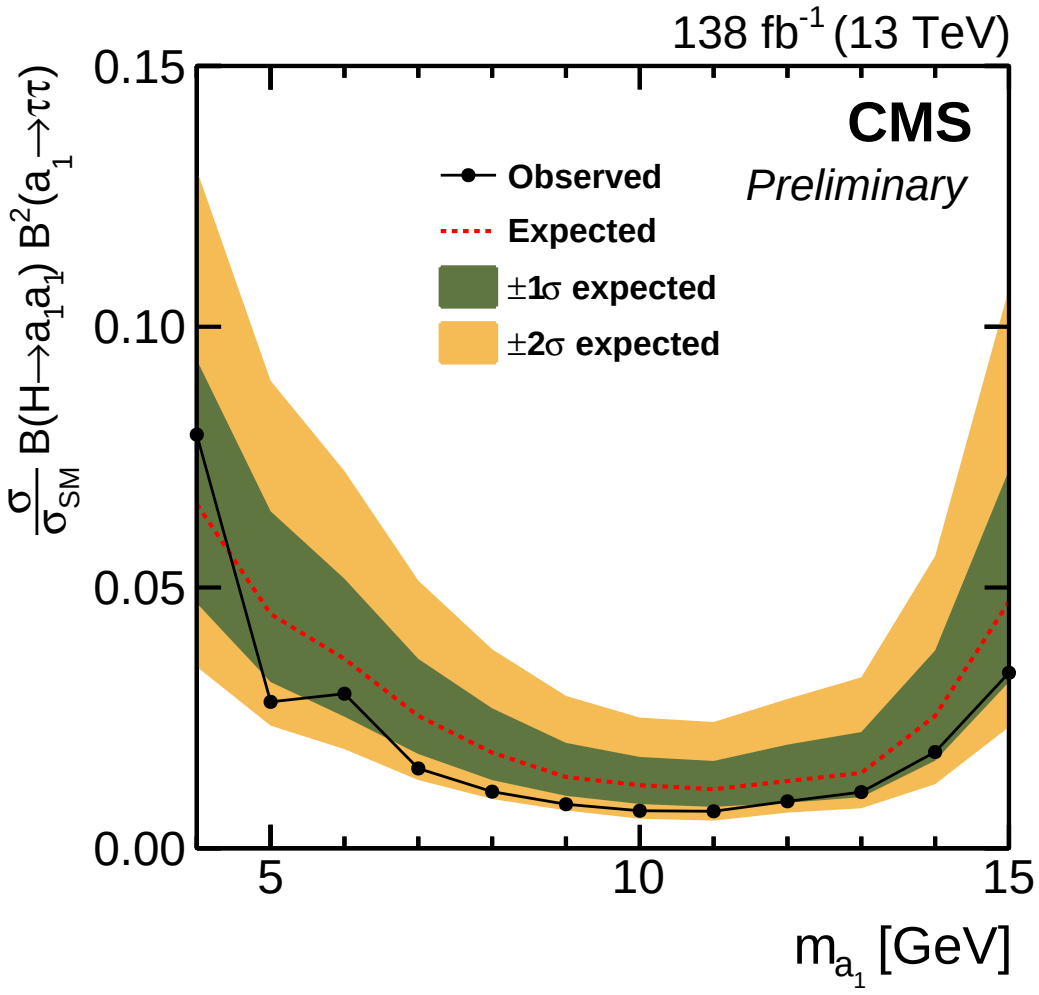

Figure 6.10: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM. These limits are plotted as a function of the $\mathrm{a_1}$ mass for Run 2. The green and yellow bands represent the 68% and 95% confidence intervals of the limits expected under the background-only hypothesis.

at lower masses. The expected limits show an improvement of 24-35%, while the observed limits are enhanced by up to 30%.

The variations in the limits across the different masses can be attributed to two distinct phenomena in the analysis. At lower m_{a_1} values, the sensitivity decreases due to the increasing similarity between the background and the signal shapes of the muon-track systems, as shown in Figs. 5.21 and 6.3. This overlap reduces the discriminatory power of the 2D (m_1, m_2) distributions. Conversely, at higher m_{a_1} values, the sensitivity drops because the average angular separation between the decay products of the a_1 boson widens, as illustrated in Fig. 5.10. As a result, the efficiency of the signal selection diminishes due to the requirement that both the muon and the track from the $\mathrm{a}_1 \to \tau_\mu \tau_{\text{1-prong}}$ or $\mathrm{a}_1 \to \mu\mu$ decay are within a $\Delta R = 0.5$ cone.

The results demonstrate a significant improvement over the previous CMS analysis performed using partial Run 2 data [197]. The limits are improved by a factor of two to four, depending on the mass hypothesis. The better limits are not solely attributed to the larger analyzed data sample. Significant improvements stem from the implementation of a veto on b-tagged jets and the more stringent criteria for the impact parameters of isolation tracks, both of which are instrumental in diminishing background levels. Additionally, the refinement of the background model, particularly through the minimization of signal contamination, has further strengthened the analysis.

6.3 Interpretation in the context of 2HDM+S Scenarios

Although the search did not yield any concrete evidence for $\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1$ decays, it can be used to constrain the parameter phase space of the different 2HDM+S models. The upper limits obtained from the experimental search are translated into constraints on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1)$ by scaling them as follows:

$$\frac{\sigma}{\sigma_{SM}}\mathcal{B}(H \to \mathrm{a}_1\mathrm{a}_1) = \frac{\left(\frac{\sigma}{\sigma_{SM}}\mathcal{B}(H \to \mathrm{a}_1\mathrm{a}_1) \cdot \mathcal{B}^2(\mathrm{a}_1 \to \tau\tau)\right)_{exp}}{\left(\mathcal{B}(\mathrm{a}_1 \to \tau\tau); m_{\mathrm{a}_1}, \tan\beta\right)^2} \tag{6.14}$$

Here, the branching fraction $\mathcal{B}(\mathrm{a}_1 \to \tau\tau)$ in the denominator is obtained from theoretical predictions and varies depending on the specific 2HDM+S model, the mass of the a_1 boson and $\tan\beta$. These values of $\mathcal{B}(\mathrm{a}_1 \to \tau\tau)$ are calculated using the decay width expressions from Ref. [215] for different points in the phase-space and provided as scans in Ref. [216].

Figures 6.11, 6.12, 6.16 and 6.16 present the observed and expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1)$ for the four types of 2HDM+S models, specifically for benchmark values of $\tan\beta$ where the $\mathrm{a}_1 \to \tau\tau$ decay exhibits a significant branching fraction. The masses considered in these plots range from 4 to 15 GeV. The experimental limits that form the basis for the numerator in Eq. 6.14, shown in Fig. 6.10, were determined for mass points generated at 1 GeV increments. For the model-dependent analysis, however, expected limits are interpolated to provide reliable estimates at intermediate masses. This interpolation is feasible due to the smooth variation of expected limits across the mass points. In contrast, the observed limits, which exhibit fluctuations at certain mass points due to statistical variations

as observed in Fig. 6.10, are maintained as discrete values. This methodical approach ensures that potential non-linearities are not overlooked by straightforward linear or spline interpolation methods and thereby guarantees that the constraints accurately reflect the experimental observations.

Additionally, 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, are established and plotted as a function of $\tan\beta$ ranging from 0.5 to 10 for selected pseudoscalar masses, as depicted in Figures 6.13 and 6.15.

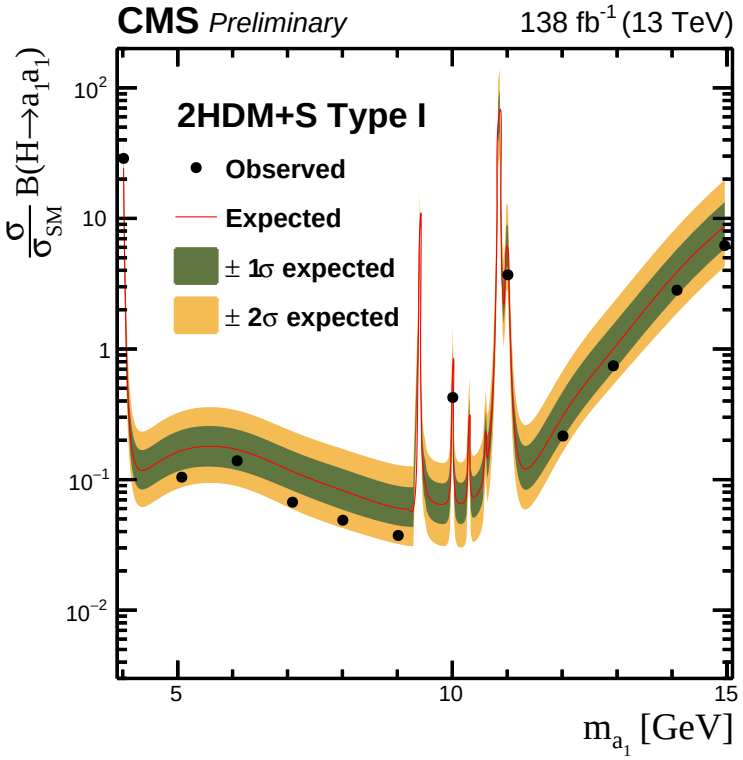

Figure 6.11: The observed and expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ for the Type I 2HDM+S model.

In the simplest 2HDM+S scenario, Type I, all fermions are coupled exclusively to a single Higgs doublet, Φ_2. Consequently, the branching ratios for the decays of $\mathrm{a_1}$ to fermions are independent of $\tan\beta$. In this model, the $\mathrm{a_1} \to \tau\tau$ decay is favored for $\mathrm{a_1}$ masses up to the b-quark pair production threshold of about 10 GeV. Beyond this threshold, $\mathrm{a_1}$ decays to b-quark pairs become dominant, driven by the larger Yukawa couplings to b-quarks. Accordingly, the upper limits set on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM}$ using the experimental results show decent exclusion capabilities for masses below the b-quark pair threshold; however, the limits weaken for higher masses, as illustrated in Fig. 6.11.

The distinct peak-like shapes observed in Fig.6.11 correspond to mass regions that coincide with quarkonium states such as η_{c} and η_{b}. In these regions, the $\mathrm{a_1}$-quarkonium mixing significantly influences the decay patterns, notably increasing the hadronic decay width through nonperturbative QCD effects. This, in turn, leads to a marked reduction in the branching fractions for decays to unbound states like $\tau\tau$. The mixing effects are particularly pro-

nounced when the couplings of a_1 to c- and b-quarks are strong. Further insights into the a_1-quarkonium mixing and its implications are detailed in Refs. [215, 217]. As a result of the mixing, the analysis fails to provide tight constraints in these mass regions.

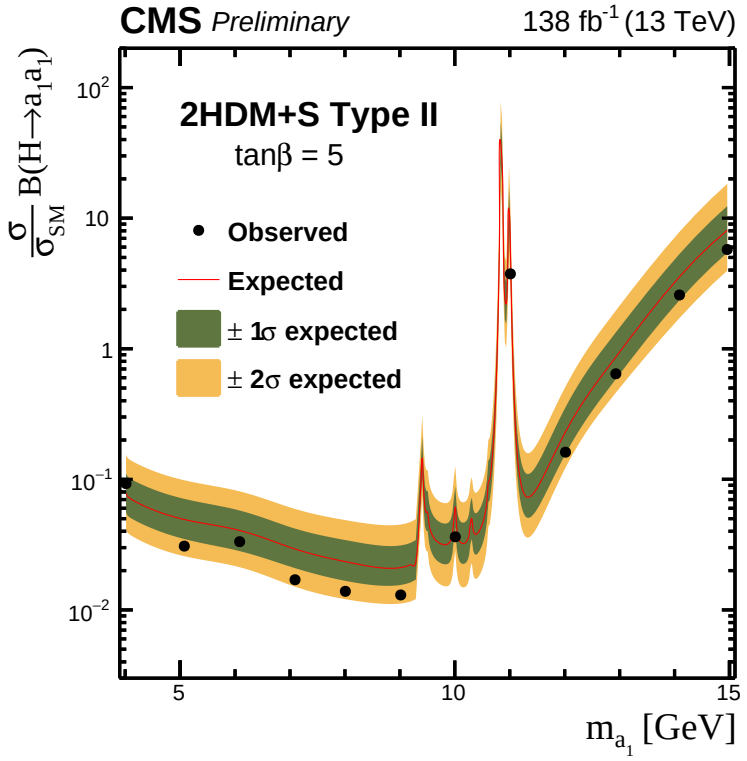

Figure 6.12: The observed and expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ for the Type II 2HDM+S model for $\tan\beta = 5$.

In the Type II 2HDM+S model, also known as the NMSSM-like version, up-type quarks exclusively couple to Φ_2, while down-type quarks and leptons interact with Φ_1. This setup results in a pronounced dependency on $\tan\beta$, which significantly influences the decay patterns of the light pseudoscalar a_1. At a high $\tan\beta$ ($\tan\beta > 1$), the interactions with down-type fermions are significantly enhanced. As a result, the $a_1 \to \tau\tau$ decay channel becomes dominant, driven by increased Yukawa couplings proportional to the fermion mass. However, once $m_{\mathrm{a_1}}$ surpasses the b-quark threshold, the decay dynamics shift; the decay $a_1 \to b\bar{b}$ becomes not only kinematically viable but also predominant. Conversely, at low $\tan\beta$ ($\tan\beta < 1$), the coupling strength to all down-type fermions diminishes, leading to a decreased branching ratio for $a_1 \to \tau\tau$. In this regime, a_1 favors decays to lighter fermions such as muons or even charm quarks, provided $m_{\mathrm{a_1}}$ is sufficiently large. The variation in decay patterns is evident in the exclusion limits presented for $\tan\beta = 5$, as shown in Fig. 6.12. For a_1 masses below 9 GeV, the analysis sets tight constraints with observed limits ranging from 0.013 to 0.09. Above 9 GeV, however, the sensitivity of the analysis is reduced.

Figure 6.13 further illustrates the exclusion limits as a function of $\tan\beta$ across several representative mass points. For a consistent $\tan\beta$ value, more stringent constraints are observed at lower masses, such as 5 and 8 GeV, while the limits become less restrictive at higher masses, such as 12 and 15 GeV. Additionally, this plot emphasizes how the branching fraction varies with $\tan\beta$. Stringent constraints are set for $\tan\beta > 1$, where the couplings to taus are enhanced. Conversely, for $\tan\beta < 1$, quark decays become dominant, leading to a suppression of $a_1 \to \tau\tau$ decays and resulting in weaker limits.

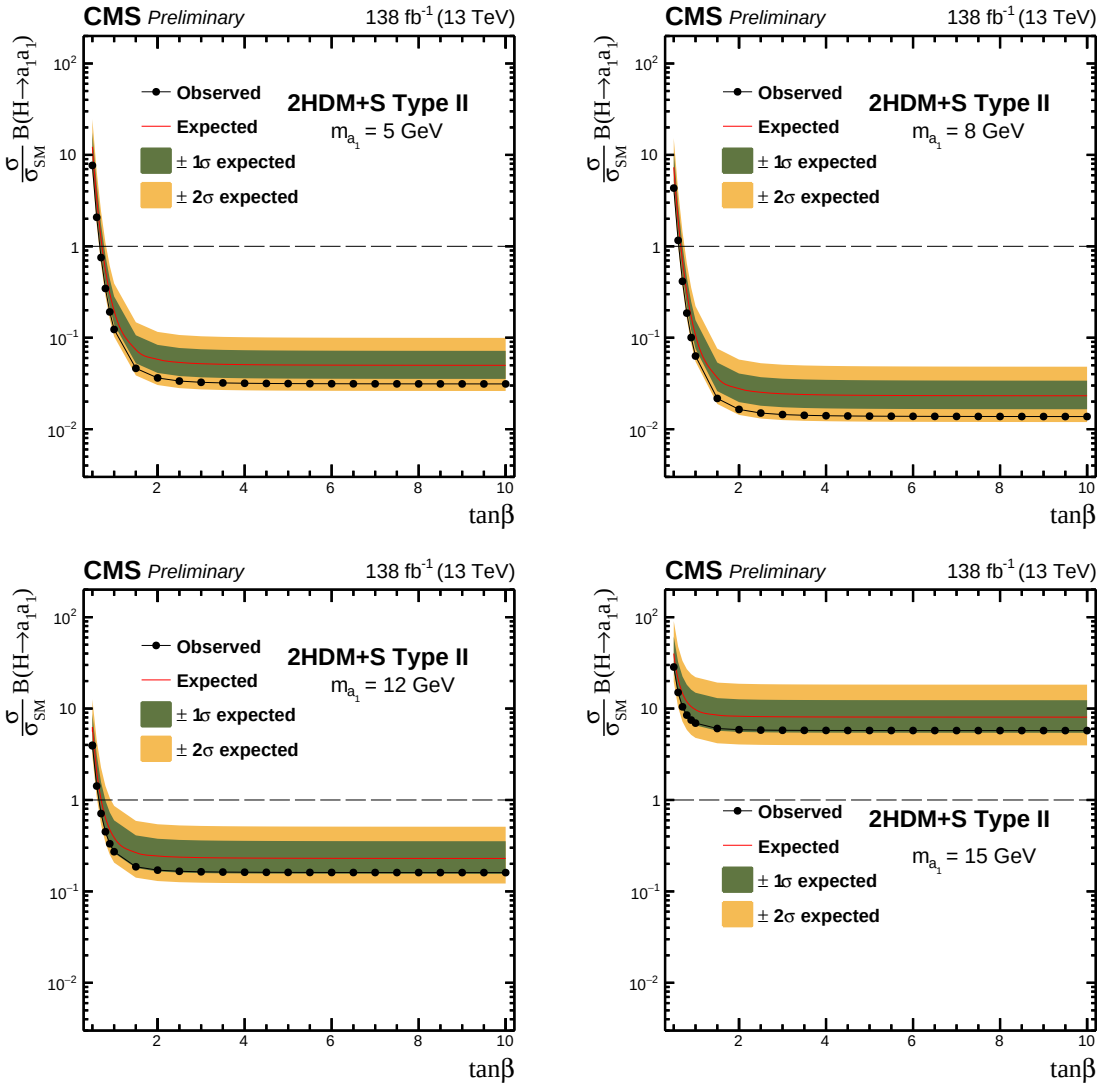

Figure 6.13: The observed and expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $\tan\beta$ for the Type II 2HDM+S model for: $m_{\mathrm{a_1}} = 5\,\mathrm{GeV}$ (upper left), $m_{\mathrm{a_1}} = 8\,\mathrm{GeV}$ (upper right), $m_{\mathrm{a_1}} = 12\,\mathrm{GeV}$ (lower left) and $m_{\mathrm{a_1}} = 15\,\mathrm{GeV}$ (lower right).

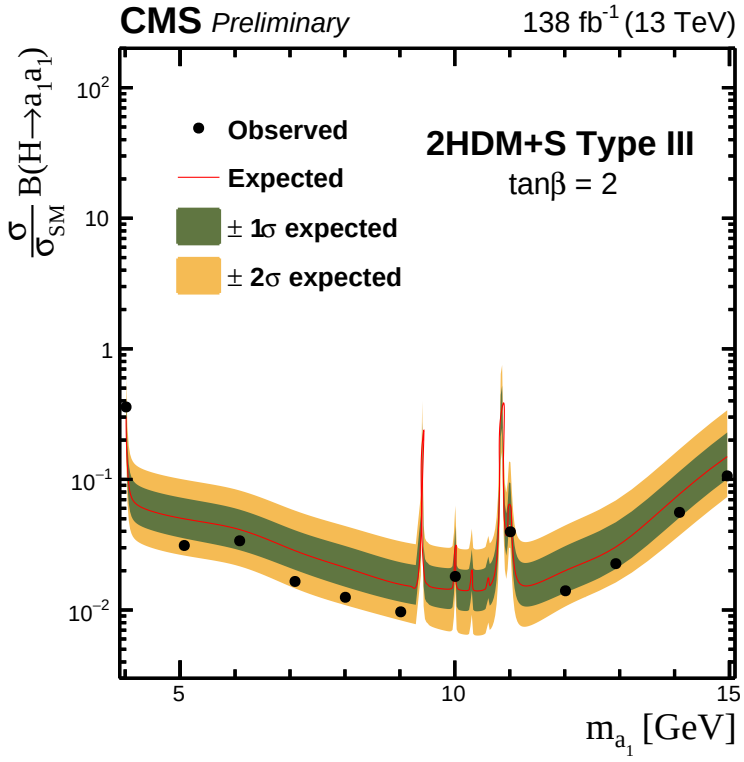

Figure 6.14: The observed and expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ for the Type III 2HDM+S model for $\tan\beta = 5$.

The Type III 2HDM+S model stands out as particularly relevant to the experimental search described in this thesis. A key characteristic of this model is the unique coupling configuration where both up-type and down-type quarks are coupled to the same Higgs doublet, Φ_2, while leptons exclusively couple to Φ_1. This setup results in a different dependency on $\tan\beta$, as compared to the Type II model. At high $\tan\beta$ values ($\tan\beta > 1$), the enhanced coupling to leptons substantially increases the branching fraction for $\mathrm{a_1} \to \tau\tau$ decays. As a result, this channel maintains its dominance throughout the entire pseudoscalar mass range, even when the masses exceed the b-quark pair threshold. This prominence is clearly evident in Fig. 6.14, which presents stringent exclusion limits for $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$ as a function of $m_{\mathrm{a_1}}$ at $\tan\beta = 2$. The observed limits here range from 0.01 at $m_{\mathrm{a_1}} = 9\,\mathrm{GeV}$ to 0.36 at $m_{\mathrm{a_1}} = 4\,\mathrm{GeV}$. The limits exclude significant portions of the model parameter space, except in regions influenced by the quarkonia effects.

At low $\tan\beta$ ($\tan\beta < 1$), the coupling to quarks, particularly to $\mathrm{c\bar{c}}$ and $\mathrm{b\bar{b}}$, dominates. This results in the experimental analysis not being able to set good limits in this region. The reduced ability to establish robust constraints in this regime is clearly demonstrated in Fig. 6.15. The figure also indicates that for $\tan\beta > 1$, the analysis achieves very stringent exclusion limits across all considered masses.

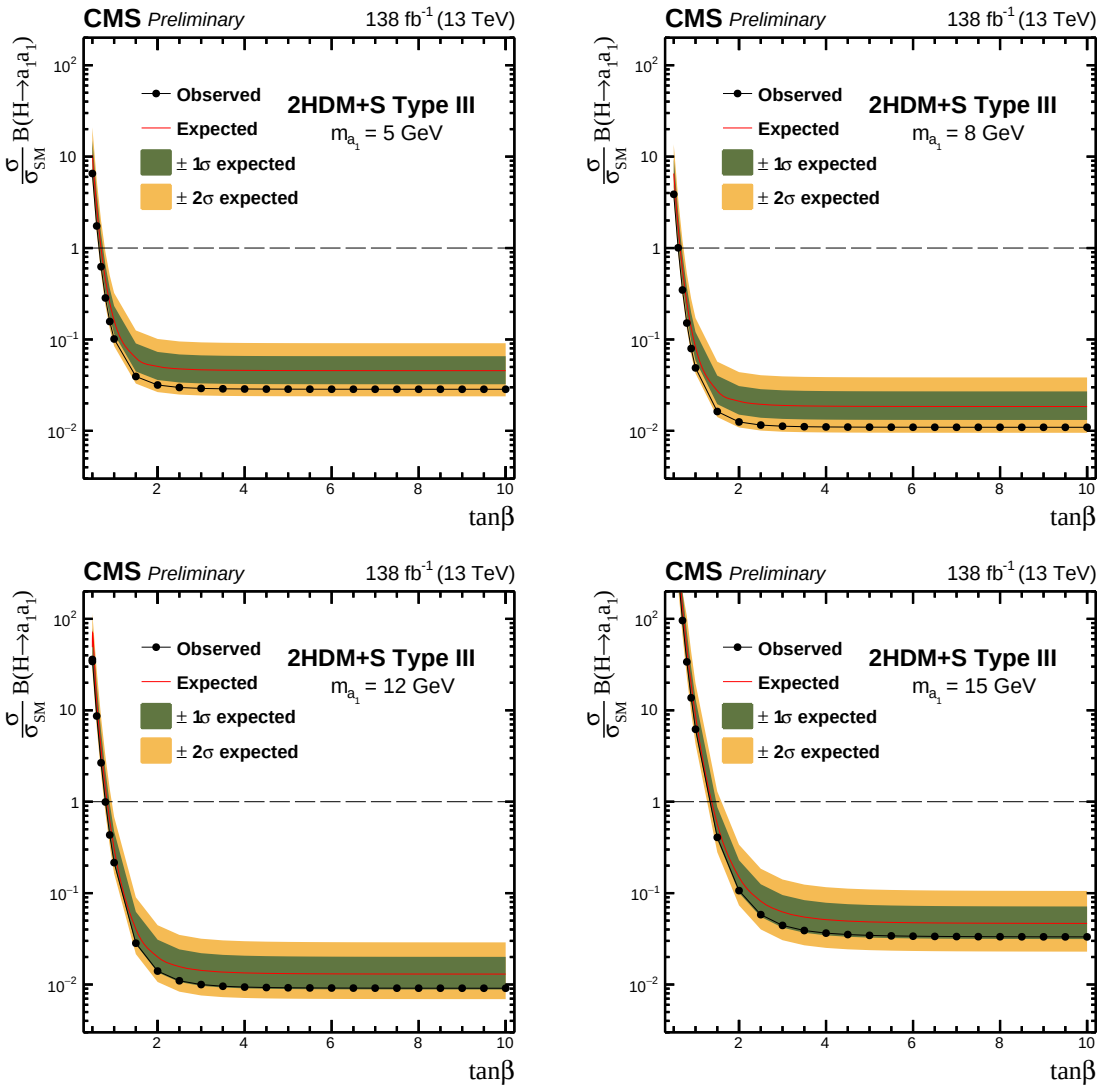

Figure 6.15: The observed and expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $\tan\beta$ for the Type III 2HDM+S model for: $m_{\mathrm{a_1}} = 5\,\mathrm{GeV}$ (upper left), $m_{\mathrm{a_1}} = 8\,\mathrm{GeV}$ (upper right), $m_{\mathrm{a_1}} = 12\,\mathrm{GeV}$ (lower left) and $m_{\mathrm{a_1}} = 15\,\mathrm{GeV}$ (lower right).

The Type IV 2HDM+S model displays markedly less sensitivity to the analysis described in this thesis. In the Type IV configuration, up-type quarks, and charged leptons are coupled to Φ_2, whereas down-type quarks exclusively interact with Φ_1. As a result, at high $\tan\beta$ values ($\tan\beta > 1$), decay modes involving down-type quarks are notably enhanced. Conversely, the $a_1 \to \tau\tau$ decay mode becomes favorable only at low $\tan\beta$ values. The exclusion limits set by the analysis for $\tan\beta = 0.5$ are illustrated in Fig. 6.16.

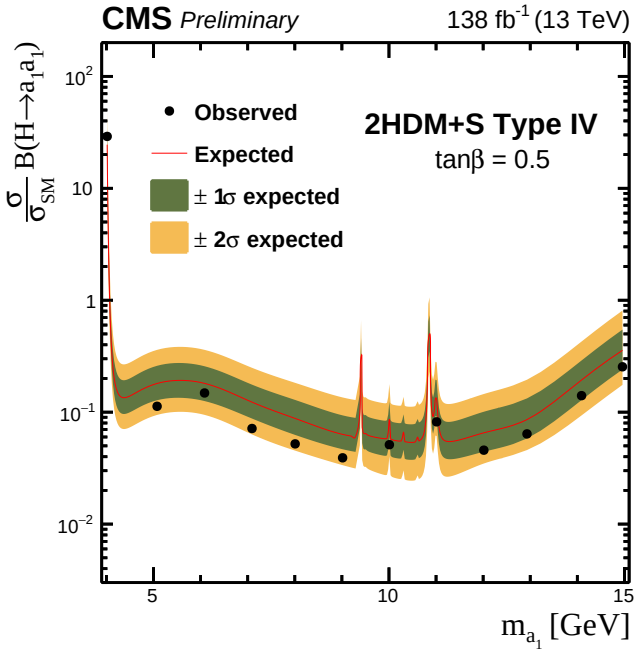

Figure 6.16: The observed and expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ for the Type IV 2HDM+S model for $\tan\beta = 5$.

Combining the mass and $\tan\beta$ dependencies into a single two-dimensional plot can provide a more comprehensive representation of how these two parameters influence the limits together. Accordingly, Figures. 6.17–6.20 show the expected limits obtained experimentally, translated into constraints on the parameters of the different 2HDM+S scenarios. The contours associated with two benchmark scenarios for the parameter $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM}$ have been added to the plot to facilitate the interpretation of the results: $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM} = 1.00$, corresponding to the case $\sigma = \sigma_{SM}$ and $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM} = 0.16$, representing the 95% CL upper limit on Higgs boson decays to BSM particles set by the CMS Collaboration using Run 2 results [9]. Conclusions similar to those derived from the one-dimensional plots can be obtained from these comprehensive two-dimensional visualizations, which also provide a more integrated view of the parameter space explored in the 2HDM+S models. The plot for Type I 2HDM+S includes the $\tan\beta$ axis for consistency; however, its uniformity along this axis clearly indicates that $\tan\beta$ plays no role in this scenario. For Type I, for $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM} = 1.00$, nearly all $m_{\mathrm{a_1}}$ values in the interval 4 to 9.2 GeV are excluded. Additionally, for $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM} = 0.16$, two main subintervals in the range of 4 to 5 GeV and 6.2 to 9.2 GeV are excluded. Some smaller subintervals at higher masses are also excluded in both cases. In contrast, the interpretations for the other model types presented in subsequent plots are more straightforward.

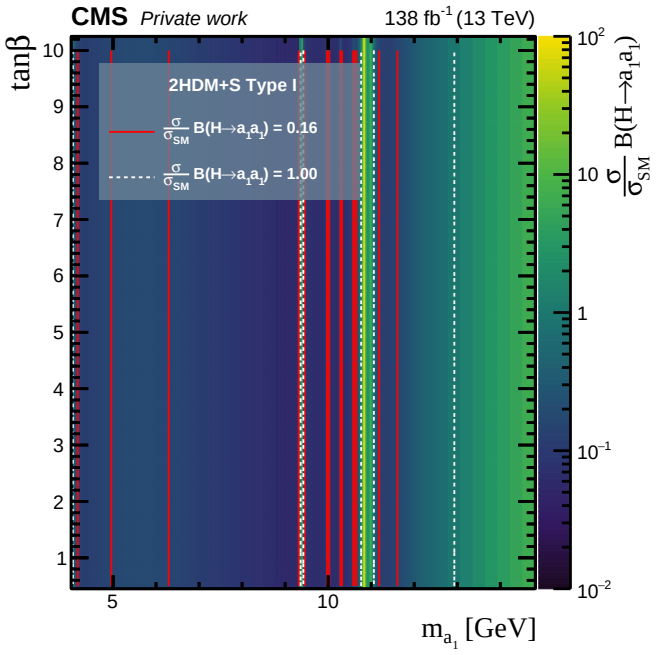

Figure 6.17: The expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ and $\tan\beta$ for the Type I 2HDM+S model. Contour lines are shown for $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1a_1})/\sigma_{SM} = 1$ and 0.16. The value of 16% corresponds to the 95% CL upper limits set on the branching fraction of the Higgs boson decays to BSM particles by CMS [9].

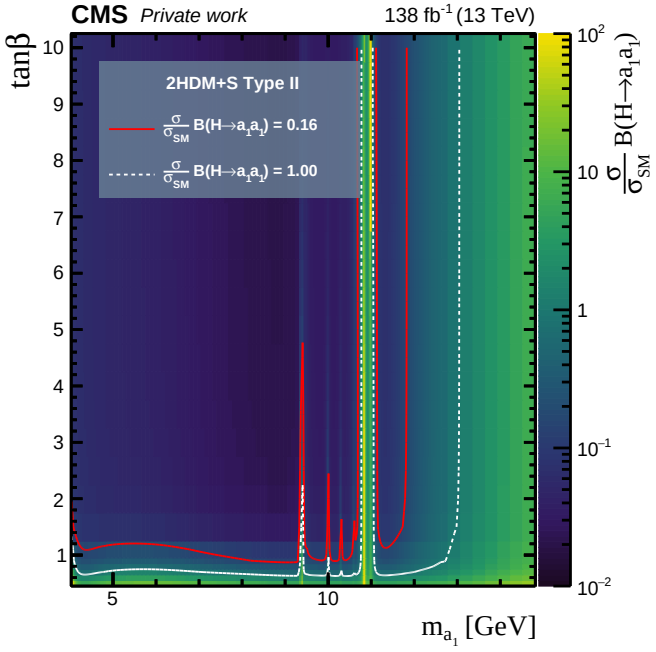

Figure 6.18: The expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ and $\tan\beta$ for the Type II 2HDM+S model. Contour lines are shown for $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1a_1})/\sigma_{SM} = 1$ and 0.16. The value of 16% corresponds to the 95% CL upper limits set on the branching fraction of the Higgs boson decays to BSM particles by CMS [9].

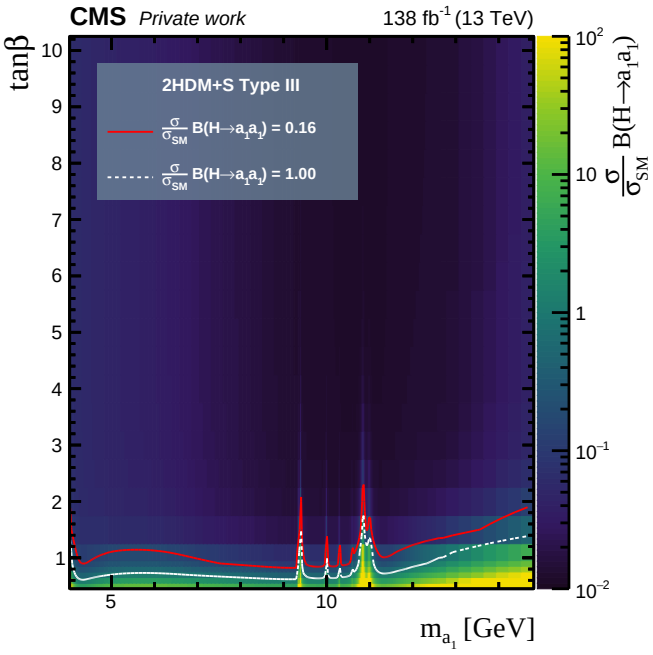

Figure 6.19: The expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ and $\tan\beta$ for the Type III 2HDM+S model. Contour lines are shown for $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM} = 1$ and 0.16. The value of 16% corresponds to the 95% CL upper limits set on the branching fraction of the Higgs boson decays to BSM particles by CMS [9].

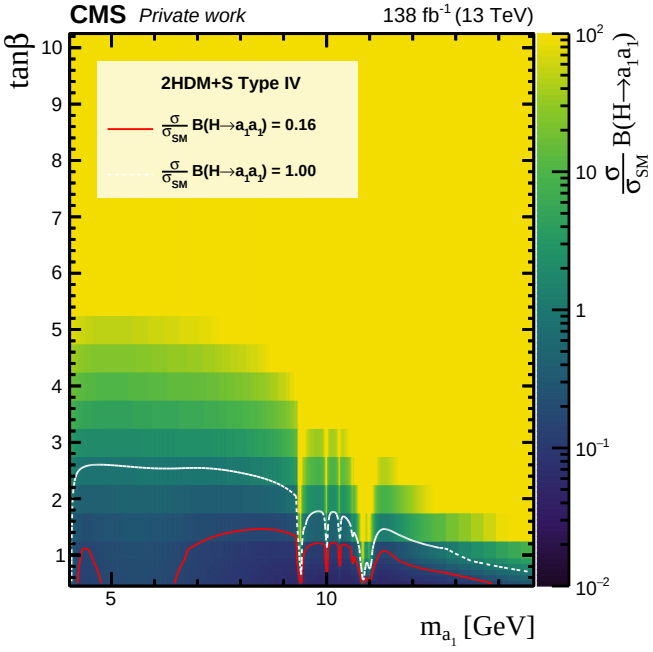

Figure 6.20: The expected 95% CL upper limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})$, relative to σ_{SM}, as a function of $m_{\mathrm{a_1}}$ and $\tan\beta$ for the Type IV 2HDM+S model. Contour lines are shown for $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})/\sigma_{SM} = 1$ and 0.16. The value of 16% corresponds to the 95% CL upper limits set on the branching fraction of the Higgs boson decays to BSM particles by CMS [9].

The analysis described in this thesis is thus able to set the most stringent constraints yet on $H \to a_1 a_1$ decays within the explored mass range for scenarios where the pseudoscalar boson's decay to leptons is enhanced. While this search in the 4τ final state provides substantial exclusion in certain areas, it is crucial to include searches targeting different final states to achieve a comprehensive understanding of the full parameter space. Both the CMS and ATLAS collaborations are intensifying their efforts through a diverse array of searches for light pseudoscalars, employing several novel techniques designed to address complex signatures. Despite these significant advances, substantial gaps persist in fully covering the parameter space. With continued advancements in experimental techniques and the huge amount of data anticipated from Run 3 and the HL-LHC phase, there is promising potential for new discoveries in the yet unexplored regions of the 2HDM+S models.

Chapter 7

Summary and Outlook

The thesis presents a search for light pseudoscalar bosons, a_1, produced from the decays of the 125 GeV Higgs boson, H, in final states with four tau leptons. To enhance the search sensitivity, the analysis also includes events from the $H \to a_1 a_1 \to 2\mu 2\tau$ process, which shares a similar topology. Such exotic decay processes are strongly motivated by theories beyond the standard model that propose an extended scalar sector. One prominent example of such a model is the Two Higgs Doublet Model with an additional scalar (2HDM+S).

To explore these theoretical predictions, the search utilizes proton-proton collision data at a center-of-mass energy of 13 TeV, collected by the CMS detector at the LHC during Run 2. The dataset corresponds to an integrated luminosity of $138\,\text{fb}^{-1}$. The analysis targets an a_1 mass range between 4 GeV and 15 GeV, as the decays of the pseudoscalar to tau leptons are theorized to be enhanced in this range within several models. The research and analysis described in this thesis were carried out independently by the author, encompassing all aspects of the study.

The decay processes considered in the analysis result in a complicated topology. Given the substantial mass difference between the Higgs boson and the pseudoscalar, the a_1 particles are highly boosted, leading to extremely collimated decay products. Consequently, the τ leptons produced from the decays of a_1 bosons cannot be reconstructed using traditional techniques, due to the overlapping fermion signatures in the final states. Therefore, to tackle this topology, a tailored strategy is employed. It involves detecting one of the tau leptons in the a_1 decay leg via its muonic decay, leveraging the CMS detector's superior muon reconstruction capabilities. The decays of the other tau lepton, which could happen leptonically or via 1-prong hadronic decays, are identified by the presence of a single, oppositely-charged track within a ΔR cone of radius 0.5 around the muon. This strategy also supports the identification of $a_1 a_1 \to (\mu\mu)(\tau_\mu \tau_{\text{1-prong}})$ decays.

In order to exploit the described topology, dedicated non-isolated double-muon triggers were set in place during the Run 2 data taking. These triggers also imposed a same-sign criteria on the two muons, which helps suppress a large amount of background coming from the background from processes such as top quark pair production ($t\bar{t}$), Drell–Yan, and diboson production, thereby enhancing the sensitivity of the search for these rare decay events. Unfortunately, due to technical constraints, this specific trigger was not available for the entire

2017 data collection period. Meticulous studies were thus conducted to identify a different trigger that would be suitable for the targeted topology. This effort successfully identified an alternate trigger, enabling the inclusion of the 2017 data, which contributed a substantial $41.5\,\mathrm{fb}^{-1}$ to the analysis.

After meeting the trigger criteria, events are subjected to a rigorous offline selection process that includes identifying two same-sign muons. Each muon is required to have only one track, meeting certain kinematic cuts, within a $\Delta R = 0.5$ around it. To enhance signal purity, the selection strategy was refined by imposing a b-jet veto on the events, to effectively eliminate backgrounds from heavy-flavor decays involving muons. Additionally, the impact parameter requirements for the tracks were tightened compared to previous analyses, further reducing backgrounds associated with heavy-flavor decays that feature slightly displaced tracks. These refinements significantly improved the signal sensitivity over the previous analysis, which utilized data collected at a center-of-mass energy of 13 TeV during 2016, corresponding to an integrated luminosity of $35.9\,\mathrm{fb}^{-1}$ [197]. The notable improvements prompted a reanalysis of the 2016 data with the updated selection criteria.

Following the offline selection, the two-dimensional (2D) distribution of the invariant masses of the two muon-track pairs is used to discriminate the signal from the background. The shapes of this 2D distribution for the signal samples were modeled from simulated samples. In order to estimate the background, which primarily consists of QCD multijet events, a comprehensive data-driven technique was employed. Control regions were defined by relaxing the strict requirement of the number of tracks around each muon. The signal is then extracted by performing a binned likelihood fit to the 2D invariant distributions.

No significant excess was observed in the data with respect to the expected standard model (SM) background. Hence, upper limits were set at 95% confidence level on the cross section times branching fraction for the $\mathrm{H} \to \mathrm{a}_1\mathrm{a}_1 \to 4\tau$ process, with respect to the SM production cross section. The statistical inference was performed individually for each analyzed year as well as for the full Run 2 combination. The full Run 2 results demonstrate a significant improvement over the previous CMS analysis performed using just 2016 data. The observed limits are improved by a factor of two to four, depending on the mass hypothesis.

The experimental results were further translated into constraints on the parameter phase space of 2HDM+S. The analysis set the tightest constraints for the Type III 2HDM+S scenario. It is able to exclude regions of the parameter space for $\tan\beta > 2$ for almost all considered masses. For all other 2HDM+S scenarios, the analysis is mainly sensitive for masses below the b-quark pair production threshold.

The results of this search have been published as a CMS Physics Analysis Summary [18]. Final preparations for journal publication are underway.

At the time of writing this thesis, Run 3 of the LHC is ongoing that is expected to deliver over twice the luminosity of Run 2, offering a substantial increase in data volume. Looking ahead, the LHC is scheduled to enter its High Luminosity Phase in 2030, which will last approximately ten years. During this phase, almost $3000\,\mathrm{fb}^{-1}$ of data is expected to be collected at a center-of-mass energy of 14 TeV. This considerable increase in data will provide an unparalleled opportunity to explore rare and exotic decay channels with unprecedented detail. To fully leverage this increased data volume, advancements in particle reconstruction

and identification techniques will be crucial. Developing algorithms capable of efficiently identifying pairs of tau leptons with overlapping decays would be particularly beneficial for the topologies similar to the one described in this analysis.

With much of the phase space of models predicting a light scalar or pseudoscalar boson remaining unexplored, there are ample opportunities for the study of exotic Higgs decays. The continued exploration of these decays is essential for enhancing the understanding of the Higgs sector.

Appendix A

Alignment of the CMS Tracker during LHC Run 3

Contents

A.1 Tracker Alignment: Introduction 160
A.1.1 Reference Systems 160
A.1.2 Track-based Alignment Algorithms 160
A.1.3 Weak Modes 162
A.1.4 Alignment Strategies 163
A.1.5 Alignment Validation 166
A.2 Tracker Alignment during LHC Run 3 167
A.2.1 Start-up Alignments 167
A.2.2 Alignment during Data Taking 170
A.2.3 Alignment for Reprocessing of 2022 and 2023 data 171
A.3 Summary 176

The innermost component of the CMS detector—the tracker—is responsible for the precise determination of the trajectory of charged particles. A detailed description of the design of the tracker is provided in Sec. 3.2.2.2. The purpose of the tracker is the precise determination of the trajectory of charged particles (tracks) from signals (hits) in the subdetector. A precise knowledge of the position, orientation, and surface deformations of each module is necessary in order to achieve optimal track parameter resolutions. However, due to mechanical tolerances during construction, as well as deformations and shifts induced by temperature variations, magnetic fields, and operational stresses, the actual positions of the tracker modules deviate from their ideal design values. Even small misalignments can lead to systematic distortions in track reconstruction, affecting the accuracy of momentum measurements and vertex reconstruction. To mitigate these effects, a dedicated calibration procedure known as tracker alignment is employed, which corrects for the position, orientation, and possible deformations of the tracker modules [100,101]. This procedure ensures that the reconstructed hit positions accurately reflect the true trajectories of charged particles, thereby preserving the precision of track parameter measurements. This chapter first provides an overview of

the methodologies and strategies followed in the tracker alignment process. The second half of this chapter highlights some of the efforts undertaken during LHC Run 3.

A.1 Tracker Alignment: Introduction

A.1.1 Reference Systems

Complementary to the global coordinate system denoted as (x, y, z) defined in Sec. 3.2.1, a local reference system (x', y', z') is defined for each module of the tracker with the origin at the geometric center of the module. The x'-axis aligns with the most precisely measured coordinate of the module, the y'-axis is orthogonal to the x-axis and lies within the module plane, pointing away from the readout electronics, and the z'-axis is perpendicular to the module plane. Angles relative to these axes are denoted as α, β and γ.

A.1.2 Track-based Alignment Algorithms

Track-based alignment algorithms are based on the principle that inaccuracies in the assumed tracker geometry during track reconstruction typically lead to increased track-hit residuals. These residuals reflect the differences between the measured and predicted hit positions, as visually demonstrated in Fig. A.1. To accurately determine the modules' positions and orientations, these algorithms minimize the sum of squares of normalized residuals across a large dataset of reconstructed tracks. The alignment process employs a least squares approach to derive alignables $\mathbf{p}$ by minimizing the following χ^2 function:

$$\chi^2(\mathbf{p},\mathbf{q}) = \sum_{j}^{\text{tracks}} \sum_{i}^{\text{measurements}} \left(\frac{m_{ij} - f_{ij}(\mathbf{p},\mathbf{q})}{\sigma_{ij}} \right), \tag{A.1}$$

where $\mathbf{q}$ represents the track parameters and the track-hit residual ($m_{ij} - f_{ij}$) is obtained by subtracting the hit prediction (f_{ij}) from the measured hit position (m_{ij}) [100]. σ_{ij} corresponds to the uncertainty in m_{ij}. The number of alignables $\mathbf{p}$ varies based on the required precision of the alignment. For aligning large mechanical structures within the pixel tracker, typically, six parameters are used for each structure's position and orientation, totaling 36 parameters. Conversely, aligning each module of the entire tracker involves eight or nine parameters per module to correct position, orientation, and surface deformations, along with extra parameters for temporal variations, culminating in $\mathcal{O}(10^6)$ alignment parameters.

An initial description of the geometry, $\mathbf{p_0}$, derived from design drawings, or mostly previous alignment results, serves as the basis for the first estimate of track parameters $\mathbf{q_{j0}}$. Since the corrections to $\mathbf{p_0}$ can be assumed to be minor, the function f_{ij} in Eq. A.1 can be linearized around the initial values, making the minimization of χ^2 analogous to solving a linear system:

$$\mathbf{Ca = b,} \tag{A.2}$$

where $\mathbf{a^T = (\Delta p, \Delta q)}$ denotes the corrections to both the geometry and the track parameters of all considered tracks. The matrix C is a correlation matrix whose elements depend on the derivatives $\partial f_{ij}/\partial p$, $\partial f_{ij}/\partial q_j$, and the measurement uncertainties σ_{ij}^m; its size corresponds

Figure A.1: Schematic illustration of how the predicted trajectory of a fitted track is influenced by the alignment of tracker modules. While local hit positions remain constant, their global positions vary based on the alignment of the modules they correspond to. Misaligned modules cause the measured hit positions to deviate from the actual trajectory, resulting in increased track-hit residuals and a distorted trajectory [218].

to the total number of track and alignment parameters. The vector b is a source term, with components that depend on $\partial f_{ij}/\partial p$, $\partial f_{ij}/\partial q_j$, σ_{ij}^m, and the measurements m_i.

In order to solve the system, the matrix $\mathbf{C}$ must be inverted. However, a full inversion is unnecessary because the matrix is sparse. Also, since the focus is primarily on the alignment parameters, it is not required to invert the sections of the matrix dealing with track parameters; it suffices to consider their correlations with the alignment parameters. Using block matrix algebra, the process is streamlined by first addressing the blocks corresponding to the track parameters, then adjusting the larger block concerning the alignment parameters and the source term. Consequently, the above equation simplifies into a system of linear equations that includes only the alignment parameters, while still accounting for all track-related correlations, as shown below:

$$\mathbf{C'\Delta p = b'}, \tag{A.3}$$

where **C'** and **b'** are obtained from **C** and **b** respectively with a significantly smaller size.

The matrix inversion in the track-based alignment process is performed within the CMS Collaboration using two independent algorithms: either globally using the MILLEPEDE-II [219] or locally using the HipPy algorithm [101]. Since the HipPy algorithm was not used to derive the results shown in the subsequent sections, the discussion henceforth will be limited to MILLEPEDE-II.

The MILLEPEDE-II algorithm consists of two main steps:

- MILLE: This program has to be integrated into the track fitting software of the specific experiment. It calculates and stores the independent residuals with errors, and derivatives of the track (local) and module (global) parameters, as per Eq. A.1. The MILLEPEDE-II employs the General Broken Lines (GBL) method for refitting trajec-

tories, starting from the parameters estimated at the first hit [220, 221]. The relevant trajectory data for MILLEPEDE-II are stored in binary MILLE files, with global derivatives included.

- PEDE: This is an experiment-independent Fortran program that constructs and solves the linear equation system described in Eq. A.3. It uses a steering file and reads track and hit data from MILLE binary files to perform local track fits and build the global matrix **C'**. This matrix is stored either fully (lower triangular part) or in a sparse format (only non-zero elements). Until 2023, for an alignment of the pixel and strip tracker at the modular level, only an approximate solution method using MINRES-QLP was possible [222, 223]. However, recently, a solution method based on LAPACK was introduced [224], enabling the exact calculation, using Cholesky factorization, for the full alignment of the pixel and strip for the first time.

A.1.3 Weak Modes

A major challenge in track-based alignment arises when the matrix **C'** in Eq. A.3 is ill-conditioned, meaning it is singular. This situation can occur when certain combinations of alignment parameters have little or no effect on the track-hit residuals, and therefore on the overall $\chi^2(\Delta\mathrm{p}, \Delta\mathrm{q})$ after linearizing the track model f_{ij}. These poorly constrained combinations are known as "weak modes." They occur when coherent shifts in the alignment parameters $\Delta\mathrm{p}$ can be compensated by corresponding changes in the track parameters $\Delta\mathrm{q}$, making their individual contributions hard to determine.

The variety of potential weak modes is influenced by several factors: the geometry and segmentation of the detector, the topology of the tracks used for alignment, and the specific alignment and track parameters. These weak modes often involve coherent module displacements such as Δr, Δz, and $\Delta\phi$ as functions of r, z, and ϕ. To effectively manage them, it is essential to incorporate additional information into Eq. A.1. This can be achieved by utilizing a diverse array of track sets with different topologies and applying various physical constraints. Particularly useful are cosmic ray muon tracks and tracks from $Z \to \mu^+\mu^-$ events.

Cosmic muon tracks, originating outside the detector, traverse both halves of the tracker, aiding in correlating movements of distant modules. This feature helps in addressing the telescope weak mode, characterized by module shifts along the z-axis proportional to their distance from the beam line, $\Delta z \propto r$. As depicted in Fig. A.2, such tracks expose this deformation by introducing a kink, which increases the distance between predicted and measured hits, thus raising the χ^2 in Eq A.1. Conversely, collision tracks can still be fitted without increasing track-hit residuals.

Tracks from muon pairs resulting from Z boson decays add an extra constraint to the alignment, as their invariant mass must align with the Z boson mass. These tracks are sensitive to the twist weak mode, where modules rotate around the beam line direction ϕ based on their longitudinal position z, with $\Delta\phi \propto z$. Figure A.3 illustrates that this twist causee a bias in the track momentum. To mitigate this, both tracks are modeled as a two-body decay rather than independently, enforcing a constraint on the reconstructed Z boson mass.

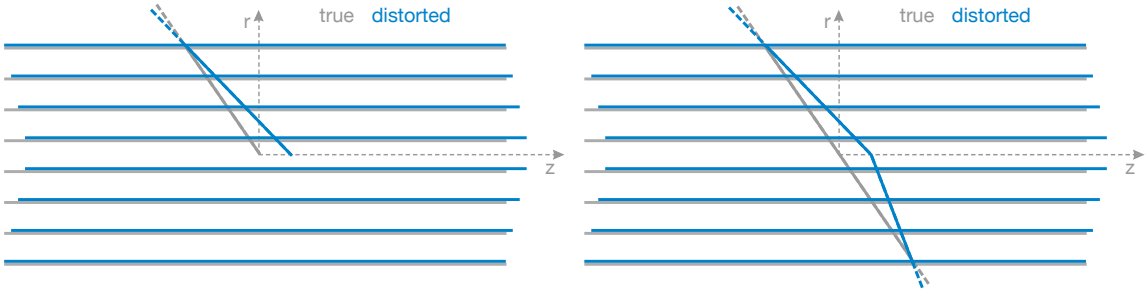

Figure A.2: Illustration of the telescope weak mode, in which $\Delta z \propto r$. The real tracker geometry (grey) is compared with the distorted one (in blue), in the rz direction. Collision tracks can be fitted in both geometries without increasing the track-hit residuals (left). This is not the case for cosmic muon tracks, for which the telescope distortion would introduce a kink in the trajectory (right) [218].

A.1.4 Alignment Strategies

The CMS tracker is organized in a hierarchical fashion, starting from large mechanical assemblies down to individual modules, as illustrated in Fig A.4. The alignment algorithm accommodates this structure by incorporating six alignment parameters for each substructure, accounting for potential translations and rotations. This hierarchical setup introduces redundant degrees of freedom because the movements of larger assemblies can be described either through their own alignment parameters or through those of their constituent parts. To manage this redundancy, the system utilizes linear equality constraints applied with Lagrange multipliers.

Particularly beneficial is the flexibility of the alignment algorithm to focus on larger assemblies, an advantage when the available track sample is insufficient for deriving alignment parameters at the module level. In such instances, aligning at the level of larger structures allows for the determination of their positions and orientations, which can then be propagated to their subcomponents. This higher-level alignment approach is crucial for effectively managing the tracker's complex mechanical structure.

Tracker alignment introduces time dependence by defining time windows called intervals of validity (IOV) to adjust for dynamic changes in the detector's geometry that occur throughout its operational life, like thermal expansion, mechanical stress, and radiation damage. These changes can influence the precision of particle tracking and impact subsequent data analysis. The time-dependent parameters are valid only in a given IOV and are determined using only tracks recorded in that time window, while the time-independent parameters use the full track sample. By continuously updating the alignment parameters to reflect these temporal variations, the alignment system maintains the high accuracy needed for particle physics experiments.

This approach, called differential alignment, is often combined with hierarchical alignment, such that the time dependency is taken into account for the high-level structures, while the positions of the single modules with respect to their high-level structure are derived using the full statistical power of the dataset.

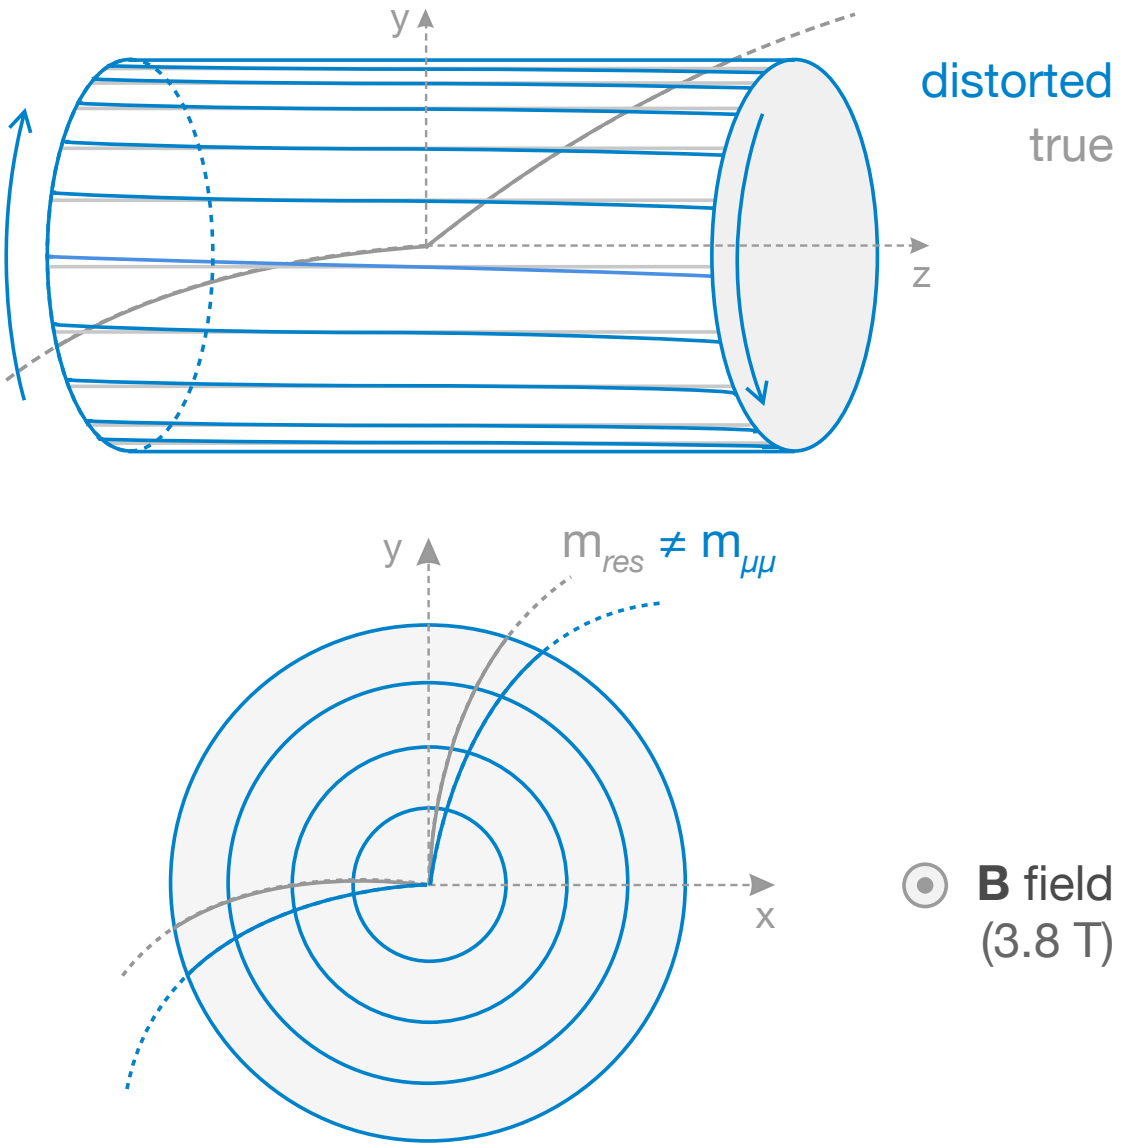

Figure A.3: Illustration of the twist weak mode, wherein $\Delta\phi \propto z$. The real tracker geometry (grey) is compared with the distorted one (in blue) (left). Tracks can be reconstructed in the distorted geometry without increasing the track-hit residuals by changing the track curvature, and hence the track momentum (right). Track pairs originating from Z boson decays are sensitive to such deformation, as their invariant mass should be compatible with that of a Z boson [218].

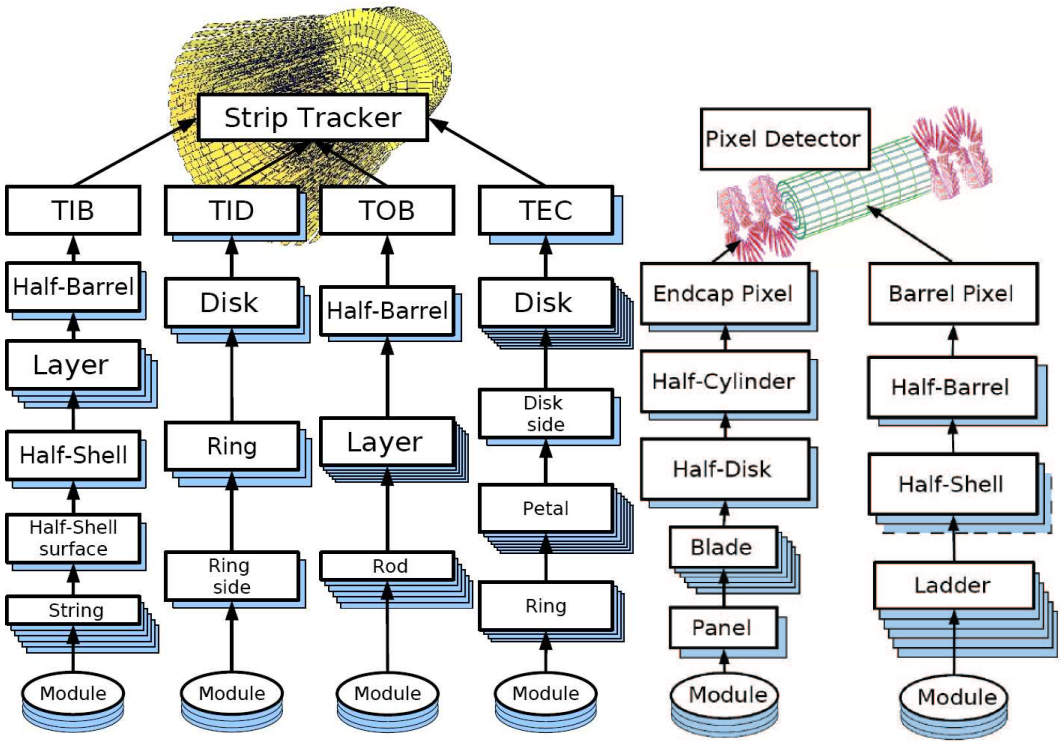

Figure A.4: Hierarchical structure of the strip tracker (left) and the pixel tracker (right) [225, 226].

A.1.4.1 Datasets Used

Different datasets are utilized for the alignment of the CMS tracker, each chosen based on its availability and characteristics. These include:

- Minimum bias events, which are collected using loose trigger conditions to capture a broad sample of inelastic proton-proton collisions. These events provide a general representation of the collision environment with minimal selection bias.
- Isolated muon events, typically originating from the decay of W bosons. The clean signature of isolated muons allows for precise tracking and makes them highly valuable for alignment purposes.
- Di-muon resonances, such as those from Z bosons or Υ mesons. These tracks are important for alignment because the invariant mass of the dimuon system is well known, providing a strong constraint for calibration.
- Cosmic ray tracks, recorded both with the CMS solenoid off (Cosmic Run at Zero Tesla, CRUZET) and on (Cosmic Run at Four Tesla, CRAFT). These tracks are especially useful for alignment as they often traverse large portions of the detector, providing coverage in regions less illuminated by collision data.

The criteria for selecting events for alignment from these datasets are detailed in Refs. [100] and [101].

A.1.5 Alignment Validation

The performance of the updated tracker geometry, following alignment corrections, is thoroughly assessed using various validation methods [100, 101]. These methods compare how certain key variables measured with the new alignment conditions stack up against an ideally aligned scenario. This section will detail some of these validation approaches.

To ensure the validation results are unbiased, it's crucial to use a set of tracks that is different from the one employed in the alignment process. This approach helps maintain the accuracy and reliability of the validation findings.

A.1.5.1 Distribution of Median Residuals

The statistical accuracy of the alignment, as well as the presence of systematic biases, is assessed using the median distribution of track–hit residuals. Each track is refitted using the proposed alignment parameters, excluding the hit being analyzed, to calculate residuals for each module, which are then used to populate a histogram. The width of this histogram, known as the distribution of the medians of residuals (DMR), gauges the precision of the alignment, while a non-zero mean suggests biases.

A.1.5.2 Primary Vertex Residual Validation

Unbiased track-vertex residuals provide a measure of the vertex reconstruction accuracy. The primary vertex (PV) position for each track is recalculated without including the track being analyzed, using the deterministic annealing clustering algorithm described in Sec. 4.2.3. In an ideal scenario with perfect alignment and calibration, the mean values of these residuals should approach zero. Random misalignments of the tracker modules influence only the distribution's resolution, broadening its width, whereas systematic misalignments shift the mean values, with the bias's direction and magnitude reflecting the nature and extent of the misalignment.

A.1.5.3 Muon Track-Split Validation

The unique topology of cosmic ray tracks, which pass through the detector vertically, is utilized for the "track-split" validation of the alignment results. In this method, cosmic ray tracks are divided at their closest point to the origin, and each segment is refitted independently using the current alignment constants. Under perfect alignment, the parameters of both track halves should match. Differences in the parameters between the two halves are analyzed; the distribution's width reflects the alignment precision, and any non-zero mean value suggests potential biases.

A.1.5.4 Dimuon validation

This validation is based on the idea that several distortions of the tracker geometry can cause a bias in the reconstructed track curvature, thus leading to a bias in the reconstructed transverse momentum of the track. This validation is most commonly performed with $Z \to \mu^+\mu^-$ events. For each selected event, the reconstructed mass is categorized based on the η and ϕ coordinates of the muons, placing each into a specific bin. A Gaussian function is then

used to fit the mass distribution within each bin, with the mean of this function representing the reconstructed mass for that bin. These bins are utilized to create profiles depicting how the invariant mass varies with η or η. The reconstructed mass is slightly below the nominal mass of 91.2 GeV, at about 90.8 GeV, mostly because of the presence of QED final-state radiation [100]. Any deviation of this reconstructed mass, whether uniform or dependent on η and η, could indicate a misalignment of the tracker.

A.2 Tracker Alignment during LHC Run 3

Building on the alignment principles introduced in the previous section, this section focuses on the achieved performance of the CMS tracker alignment. A brief summary of the performance of the tracker alignment during Run2 is provided in Appendix B. The present section details some of the major efforts undertaken during LHC Run 3 to ensure good alignment of the CMS Tracker. The author of this thesis has contributed substantially to most of these results, which has resulted in the publications in Refs. [12–17]. This section is organized by the different phases of alignment executed throughout the year, each tailored to specific objectives.

A.2.1 Start-up Alignments

Start-up alignments are conducted prior to the commencement of proton-proton collisions at the start of each data-taking year. During long shutdowns or year-end technical stops, mechanical stresses, thermal cycling, and interventions such as adjustments or maintenance activities can cause shifts or unintentional movements in the positions of the tracker modules. To address this, cosmic ray data and collision data at a lower center-of-mass energy of 900 GeV, collected during the commissioning phase, are utilized for alignment purposes. The goal is to establish a good baseline geometry that provides crucial initial calibration of the detector before extensive data taking begins.

The start of Run 3 data taking presented significant challenges for tracker alignment. As detailed in Sec. 3.2.2.2, the pixel detector underwent extensive refurbishments and upgrades; notably, layer 1 of the barrel pixel was fully replaced. Following its reinstallation, substantial shifts in alignment were anticipated. Multiple iterations of alignment adjustments were thus carried out to correct for the movements.

The initial iterations were carried out using CRUZET data collected during 2021. The first round of alignment utilized 120 000 cosmic ray tracks recorded at 0 T. The alignment corrections were derived at the level of half-barrels and half-cylinders for the pixel and strip detectors, respectively. This alignment was further refined using 1.5M cosmic ray tracks, allowing an increase in the granularity of the barrel pixel detector to the level of single ladders. The refined alignment played a huge role in recovering from the misalignments caused by the extraction and reinsertion of the pixel detector. Figure A.5 showcases the DMRs for the two alignment iterations performed with the CRUZET data, compared to the last tracker geometry derived during Run 2.

Following this, CRAFT data was collected after the magnetic field was ramped up to 3.8 T. An alignment geometry was derived, using 765k cosmic ray tracks, at the level of single modules for the barrel pixel, half-cylinders for the forward pixel, and half-barrels or half-

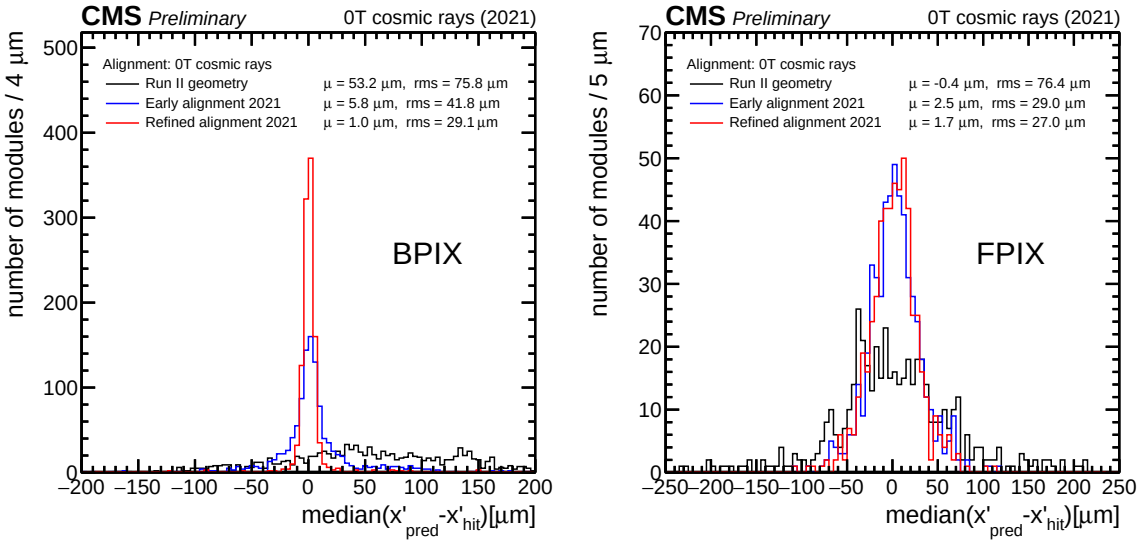

Figure A.5: Distribution of median residuals in barrel pixel (BPIX) (left) and forward pixel (FPIX) modules (right), along the local-x (x') direction. The startup geometry (black) is compared with the first CRUZET iteration (blue) and the refined alignment (red). Significant improvement is achieved with the alignment campaigns in 2021 in both BPIX and FPIX. The enhanced performance in the BPIX for the refined alignment is achieved by the increased granularity used [227].

cylinders for the different strip subdetectors. The next phase of the LHC commissioning involved the first collisions of Run 3 at a center-of-mass energy of 900 GeV. Following the start of the collisions, another geometry was derived using both cosmic and collision data, wherein the tracker was aligned at the level of single modules for the pixel detector and the different strip subdetectors. This alignment utilized about 3.6M cosmic ray tracks and 255.2M collision tracks. The cosmic rays used included cosmic rays recorded, both before and during collisions, at 0 T magnetic field, as well as 3.8 T. The DMRs for the described alignments are shown in Fig. A.6. The results of the MTS validation are also showcased in Fig. A.7. Considerable improvements were seen in the alignment performed with collisions + cosmic data. This can be mostly attributed to the finer granularity of the tracker considered for the alignment.

Start-up alignments were also performed during the year 2023 using cosmic ray tracks collected at 3.8 T. Unfortunately, the cosmic data collection was rather short this year, and only a total of 186.9K tracks could be used for deriving the alignment. With these limited statistics, alignment corrections were derived at the level of ladders for the barrel pixel detector, half-cylinders for the forward pixel detector, and at the level of half-barrels and half-cylinders for the strip partitions. Further refinements were brought about by utilizing collisions at a center-of-mass energy of 900 GeV, followed by collisions at 13.6 TeV. The alignment with 13.6 TeV collision data utilized 133K cosmic ray tracks and 16.1M collision tracks, out of which 410.8K tracks originate from pp collision events with $Z \to \mu\mu$ decays. Alignment corrections were derived at the level of ladders and panels for the pixel detector, while the strip detector was kept fixed. Figure A.8 provides a comparison of the performance of these alignments in terms of their DMR distributions and PV residual distributions. Considerable

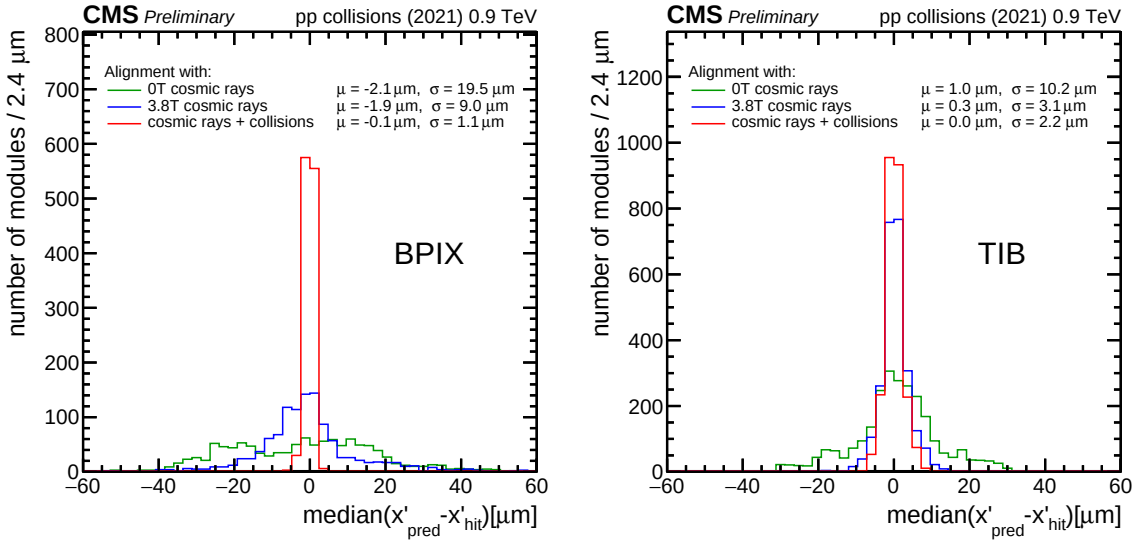

Figure A.6: Distribution of median residuals in barrel pixel (BPIX) (left) and tracker inner barrel (TID) modules (right), along the local-x (x') direction. The performance of the alignment geometries derived with CRUZET (green), CRAFT (blue) and cosmic rays + collisions at 900 GeV (red) is compared [13].

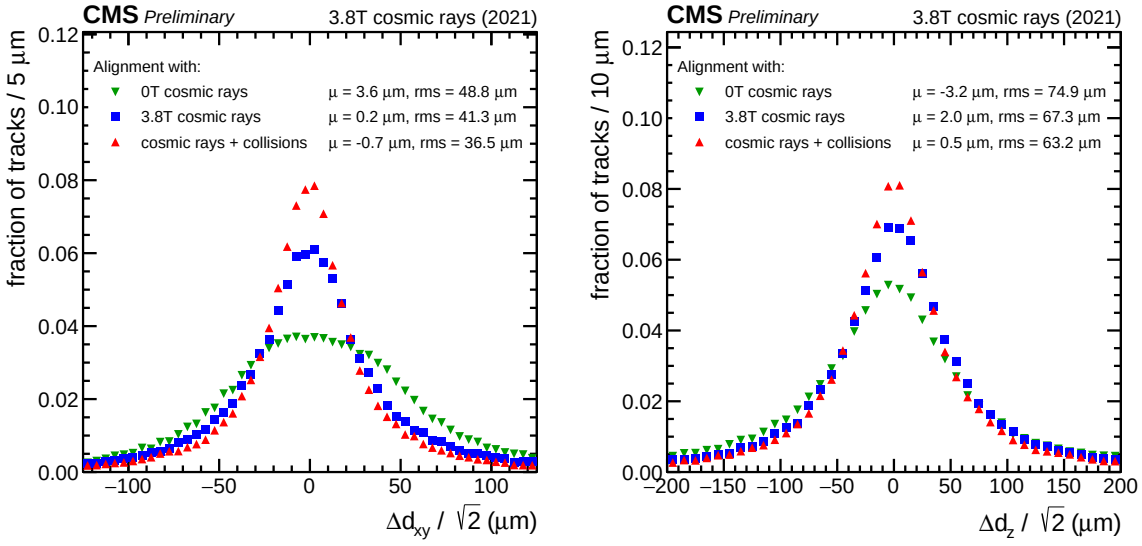

Figure A.7: Difference in the track impact parameters in the transverse plane (d_{xy}) (left) and the longitudinal direction (d_z) (right) when cosmic ray tracks passing through the pixel detector are split into two halves at their point of closest approach to the interaction region. The performance of the alignment geometries derived with CRUZET (green), CRAFT (blue) and cosmic rays + collisions at 900 GeV (red) is compared [13].

improvement is observed with the alignment using cosmic + 13.6 TeV collision data, thus providing a good initial calibration for the early 2023 collision data.

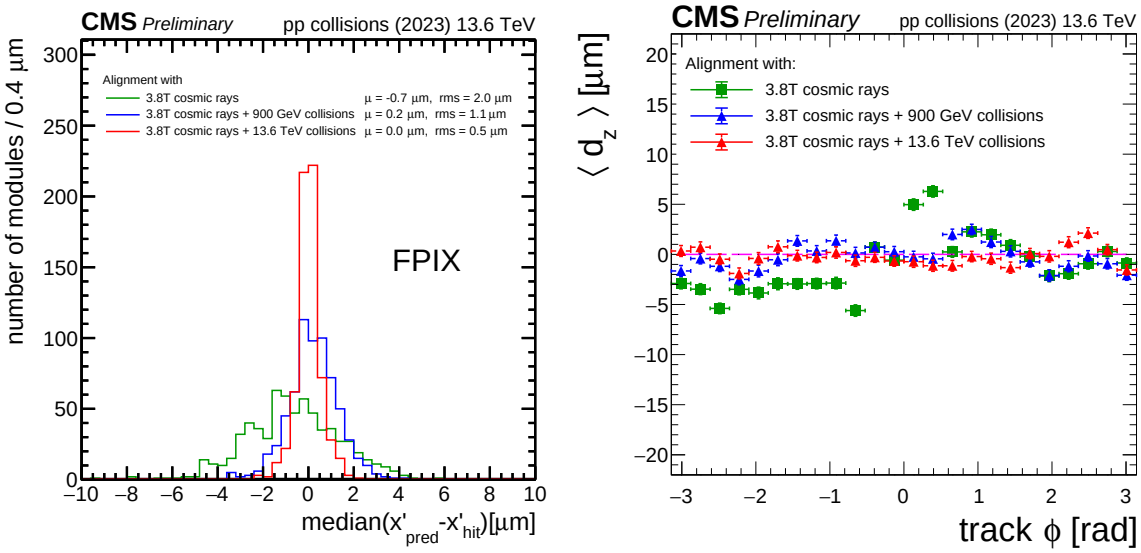

Figure A.8: The DMR distribution in the local x coordinate of the FPIX detector (left) and the PV validation for the longitudinal impact parameter (d_z) relative to the track ϕ (right). The first 2023 alignment, derived with cosmic rays (green), is compared to the next iteration, employing cosmic rays and collision tracks recorded at a center-of-mass of 900 GeV (blue), followed by an alignment with cosmic rays and collision tracks recorded at a center-of-mass of 13.6 TeV (red). The latest campaign delivers a considerable improvement [16].

A.2.2 Alignment during Data Taking

During data taking, components of the pixel detector may shift due to changes in magnetic fields or temperature variations. To address these shifts, an automated alignment procedure was introduced in Run 2, enabling fast updates to alignment parameters within 48 hours. This system, called the low granularity prompt calibration loop (LG PCL), aligns the pixel detector at the level of half barrels and half cylinders. However, it was observed that this workflow could not fully mitigate the effects of radiation damage, as the granularity of the alignment parameters used was too coarse to account for these changes. Therefore, in mid-2022, during Run 3, a new automated workflow, referred to as the high granularity prompt calibration loop (HG PCL), was implemented that aligns individual ladders and panels of the pixel detector with much finer granularity than the previous system, managing over 5,000 parameters. More details about this workflow are available in Appendix C.

Since mid-2022, this workflow, exploiting minimum bias data, has been utilized to frequently update the alignment conditions during data-taking. The high-granularity automated workflow has been successful in mitigating systematic biases caused by radiation damage. However, the lack of variety in the dataset used makes it not so good at handling weak modes. Efforts were put into modifying this workflow to include dimuon events, and an upgraded version of the workflow has also been in place since mid-2024.

Apart from the automatic workflows, manual updates are also performed from time to time

in order to provide a good starting geometry for the automated workflows.

A.2.3 Alignment for Reprocessing of 2022 and 2023 data

Data and Monte Carlo simulations are frequently reprocessed using the latest calibrations to enhance physics performance. The derivation of alignment constants for physics analysis demands large samples to precisely determine the positions, orientations, and surface deformations of all sensors in both the pixel and strip detectors. Hence, they are usually performed in the middle or at the end of the year, once a lot of statistics have been collected.

At the end of 2022 and 2023, CMS made the decision to reprocess the data collected during these years. Taking advantage of this, the tracker alignment in these years was optimized to improve the precision of physics data reprocessing. The alignment was performed considering data sets of different varieties, including minimum bias events, cosmics, events with isolated muons, and dimuon events. Each data taking year was aligned separately. To maximize the statistical power and prevent systematic distortions, data collected over an entire year were aggregated to perform the alignment fit. Temporal variations within the year were addressed by employing a hierarchical approach in the alignment fit. Both the pixel and the strip detectors were aligned at the modular level. Notably, this set of alignment parameters was derived using an exact solution method, described in Sec. A.1.2, for the first time, marking a significant advancement in the precision achieved.

The Intervals of Validity (IOVs) for each set of alignment constants were established based on various factors, including magnet cycles, modifications in hit reconstruction due to local calibration changes, voltage adjustments, and sensor aging from irradiation. Specifically, the barrel pixel modules were aligned with the smaller IOVs of the automated workflow.

However, defining IOV boundaries is only one component of achieving the required alignment precision. The alignment strategy must be tailored to the desired accuracy and data availability by adjusting the number and type of alignables, thus controlling the degrees of freedom in the alignment. Since both the alignment computation and subsequent physics validation are resource-intensive—each matrix inversion can take up to a day, and validation can span several days—the number of configurations that can be tested in parallel is inherently limited.

The derivation of the final set of alignment constants for one data-taking year thus takes several weeks and involves the efforts of multiple people.

Having outlined the workflow for deriving alignment constants for physics analyses, the subsequent discussion now focuses on evaluating their performance compared to the alignment during data taking.

Prior to mid-2022, the LG PCL performed automatic updates of the alignment constants during data taking at the granularity of half-barrel and half-cylinder movements of the pixel detector. Following this, it was replaced by the HG PCL, which aligns the pixel detector at a finer level of ladders and panels. The HG PCL was predominantly active for the whole of 2023. In both 2022 and 2023, intermittent manual updates of conditions were also done at the finest granularity of the pixel detector.

Figure A.9 shows the DMRs along the local x coordinates of the modules for the BPIX and the TID, evaluated using single muon events. The figure compares the performance of the

alignment during data taking and the alignment for reprocessing. The distributions shown are averaged over all processed runs of 2022, after scaling them with the corresponding luminosity for each run. Pixel detector modules with fewer than 100 hits per run over a period of data taking corresponding to an integrated luminosity of more than $2\,\mathrm{fb}^{-1}$ were excluded from this study. Compared to the alignment during data taking, considerable improvements were observed in both the pixel and strip detector modules. The mean value of the distributions of median residuals was shifted closer to zero, and the width was also reduced. The same was observed for the 2023 alignment as well. Improvements in the alignment for reprocessing can be attributed to the finer granularity of the alignment.

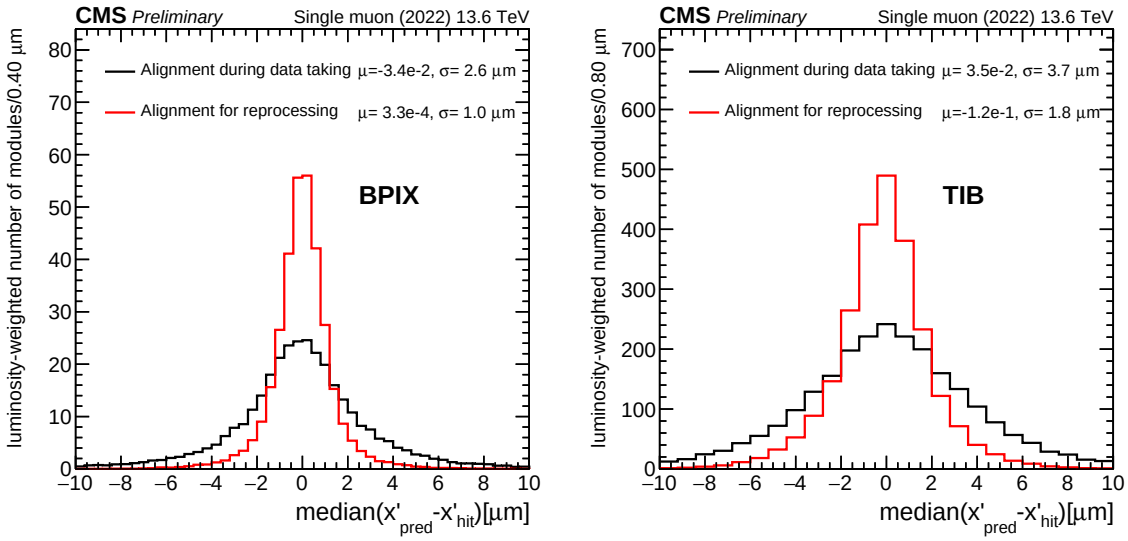

Figure A.9: Distribution of the median residuals in the BPIX (left) and TIB (right) in the local x (x') direction. The black line corresponds to results obtained with the alignment constants used during data taking, and the red line shows results with the alignment constants used for the reprocessing of 2022 data. Distributions shown here are averaged over all processed runs of 2022, after scaling them with the corresponding luminosity for each run. Pixel detector modules with fewer than 100 hits per run over a period of data taking corresponding to an integrated luminosity of more than $2\,\mathrm{fb}^{-1}$ were excluded from this study. Improvements of the mean and resolution in the alignment for reprocessing can be attributed to the finer granularity of the alignment [17].

Figure A.10 shows the mean track-vertex impact parameter in the transverse plane d_{xy} and in the longitudinal direction d_z as a function of ϕ (top) and η (bottom). The alignment for reprocessing achieves an excellent performance, with a flat mean close to zero. Modulations and biases observed in these distributions with the alignment constants used during data-taking are mitigated through alignment improvements in reprocessing. The larger bias in d_z vs. η is attributed to the sole reliance on minimum bias events in the automated alignment. Incorporating $\mathrm{Z} \to \mu\mu$ events with mass and vertex constraints in the reprocessing alignment significantly reduces this bias.

As described in Sec. A.1.5.4, muonic decays of Z bosons provide a standard reference that can be used for validating the aligned geometry. Figure A.11 shows the mass of the Z candidates as a function of the difference in η between the positive and negative muon. A sigmoid shape,

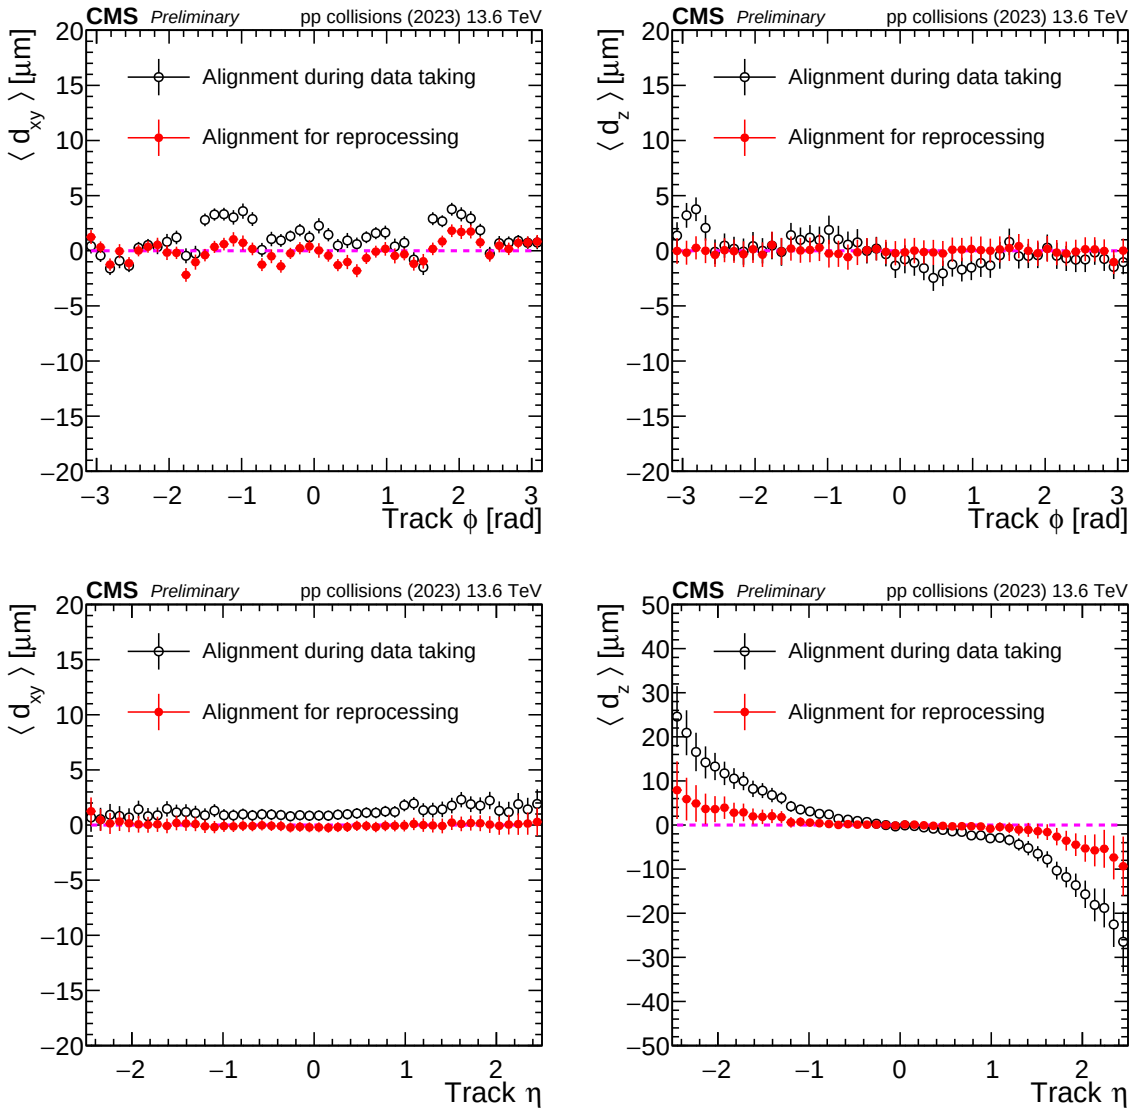

Figure A.10: Mean track-vertex impact parameter in the transverse plane d_{xy} (left) and in the longitudinal direction d_z (right) as a function of ϕ (top) and η (bottom). The impact parameters are obtained by recalculating the vertex position after removing the track being studied from it and considering the impact parameter of this removed track. The black points correspond to results obtained with the alignment constants used during data-taking, and the red points show results with the alignment constants to be used for reprocessing of 2023 data. Distributions shown here are averaged over all processed runs of 2023, after scaling them with the corresponding luminosity for each run [17].

typical of the twist weak mode, is observed in the alignment during data taking, which can again be attributed to the lack of dataset variety used in the workflow. The alignment for reprocessing shows an improvement in the uniformity of the reconstructed Z mass.

Additional systematic distortions have been mitigated through alignment improvements in reprocessing. For example, the bias in the reconstructed Z mass as a function of the azimuthal

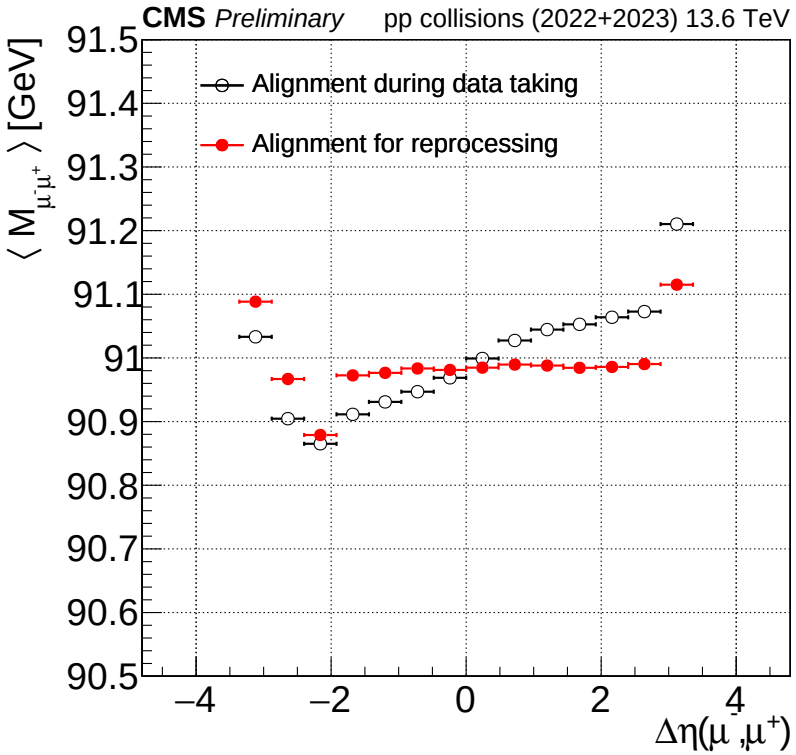

Figure A.11: Reconstructed Z boson mass as a function of the difference between the η of the negatively and positively charged muons, calculated from the full sample of dimuon events in the years 2022 and 2023. The black points correspond to the results with the alignment constants used during data taking, the red points show the results with the alignment constants to be used for the data reprocessing of 2022 and 2023. The alignment for reprocessing shows an improvement in the uniformity of the reconstructed mass [17].

angle ϕ follows a first-order periodic distortion, which can be effectively modeled using a cosine function. Consequently, the amplitude of the fitted cosine function serves as a reliable metric for quantifying the magnitude of the bias in the reconstructed dimuon mass. Figures. A.12 and A.13 illustrate the $Z \to \mu\mu$ mass as a function of the azimuthal angle ϕ of a positively charged muon. Plots are shown for different scenarios, where each muon is within a particular η region. It can be observed that the alignment during data-taking introduced a significant bias in these distributions, which was recovered by the alignment for reprocessing, thus offering ultimate physics precision.

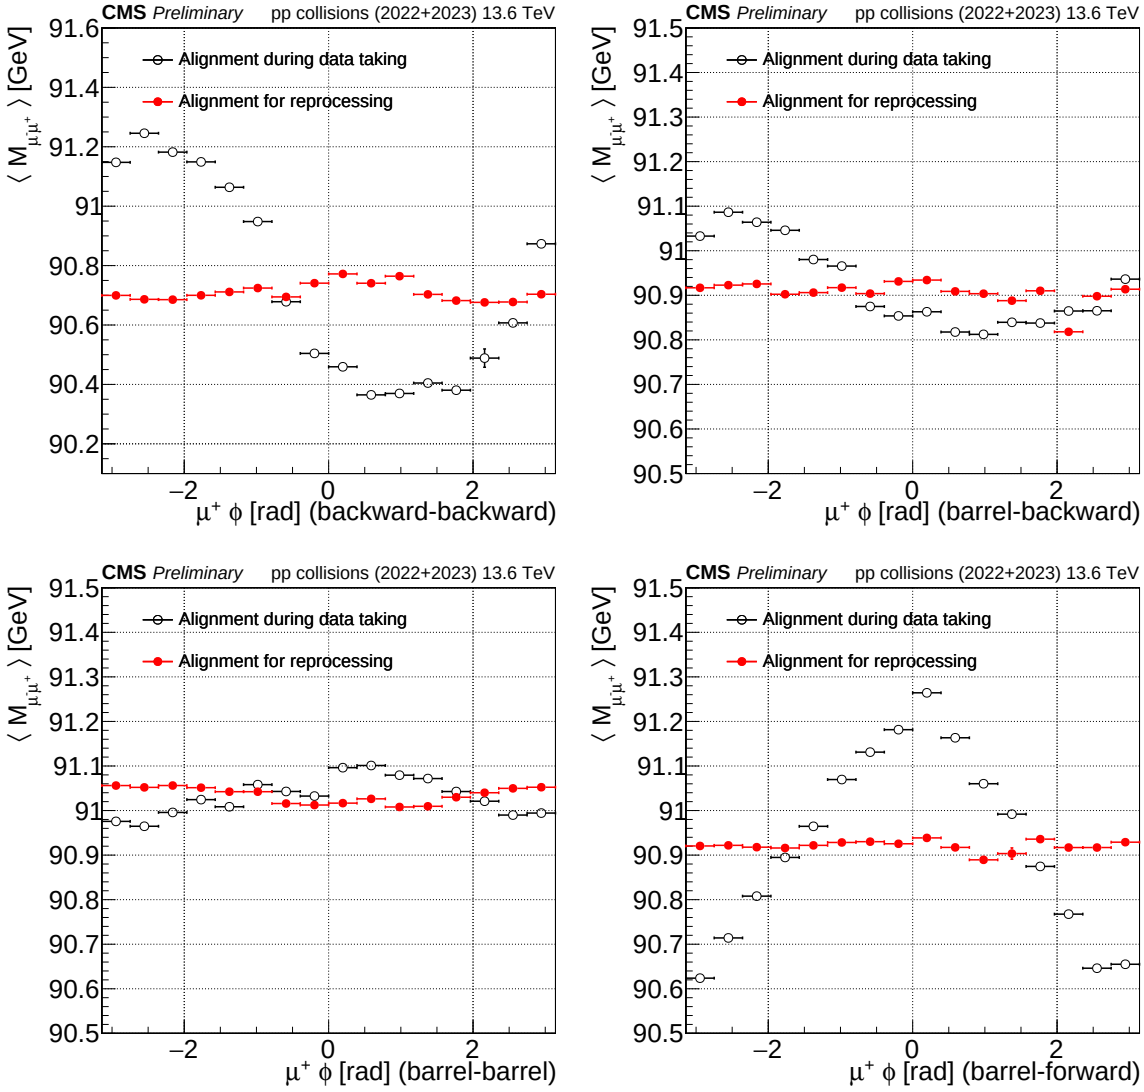

Figure A.12: Reconstructed Z boson mass as a function of the azimuthal angle ϕ for a positively charged muon, calculated from the full sample of dimuon events in the years 2022 and 2023. The plot is shown for the η region where: (top left) both muons are backward ($\eta < -1.5$; (top right) one muon is within the barrel ($\eta \leq 1.5$), and the other muon is backward ($\eta \leq -1.5$); (bottom left) both muons are within the barrel ($\eta \leq 1.5$); and (bottom right) one muon is within the barrel ($\eta \leq 1.5$) and the other muon is forward ($\eta > 1.5$).

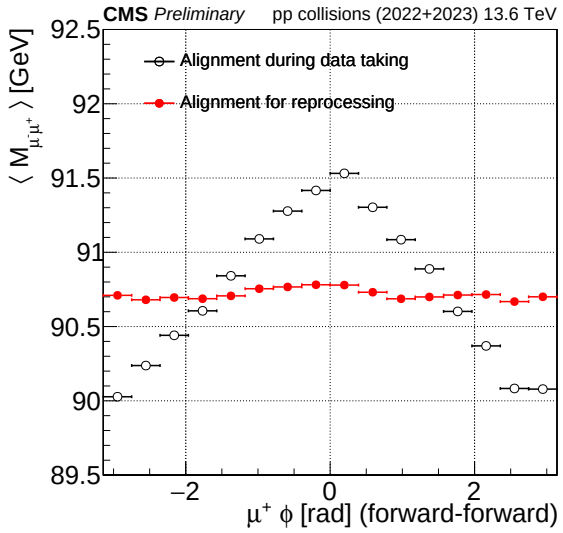

Figure A.13: Reconstructed Z boson mass as a function of the azimuthal angle ϕ for a positively charged muon, calculated from the full sample of dimuon events in the years 2022 and 2023. The plot is shown for the η region where both muons are forward ($\eta > 1.5$).

A.3 Summary

The alignment of the CMS tracker is a key part of the detector's calibration. It involves determining the precise positions and orientations of the silicon modules to ensure accurate reconstruction of particle tracks and interaction vertices. CMS uses a track-based alignment approach, described in Sec. A.1.2, which relies on residuals from reconstructed tracks to constrain module positions. The MillePede-II algorithm is used to efficiently solve the large linear system that arises in this process. To verify the quality of the resulting alignment constants, a series of dedicated validation procedures is applied. The task is challenging both computationally—due to the large number of parameters—and methodologically, since the optimal choice of alignment granularity, track selection, and statistical input is not known a priori and must be determined empirically. These general principles guided the alignment efforts during Run 3, which posed new challenges due to changes in the detector and harsher running conditions. The beginning of Run 3 data taking required substantial alignment efforts to correct for possible misalignments that were introduced by the replacement of the innermost barrel pixel layer. This was addressed through multiple dedicated campaigns using cosmic ray data and early proton-proton collisions at 900 GeV and 13.6 TeV, resulting in a well-calibrated initial geometry. As detailed in 3.1.4, Run 3 introduced higher luminosities and pileup, leading to increased radiation damage, which, in turn, placed greater demands on alignment precision. To address these challenges, new automated workflows were developed to provide continuous alignment updates, properly cope with radiation damage, and enhance the quality of data taking. Large-scale reprocessing campaigns for the 2022 and 2023 datasets achieved alignment at the level of ultimate precision. The work carried out by the author during their PhD was central to all these developments. While continuous efforts are underway to cope with detector ageing, the overall performance achieved so far in Run 3 has been excellent.

Appendix B

Tracker Alignment in Run 2 and Commissioning for Run 3

This section provides an overview of the alignment strategies employed for the CMS Tracker during LHC Run 2. It includes a detailed discussion on the performance of the tracker alignment conditions that were derived, highlighting the precision achieved through the legacy reprocessing efforts. Additionally, this section dives into the preparations made by CMS towards tracker alignment during Run 3, particularly focusing on the initial alignment conducted after the LHC Long Shutdown 2 using cosmic ray muons and collision data at $\sqrt{s} = 900\,\text{GeV}$. The content is presented as originally published in [228].

Performance of CMS Tracker Alignment during LHC Run 2 and Commissioning towards Run 3

Lakshmi Pramod[1] for the CMS Collaboration
[1]Deutsches Elektronen-Synchrotron (DESY), Notkestraße 85, 22607 Hamburg, Germany
E-mail: lakshmi.priya.sreelatha.pramod@cern.ch

Abstract. The CMS experiment's innermost detector, the tracker, is designed to accurately measure the momentum of charged particles and reconstruct the primary vertices. As the operating conditions during data-taking change frequently, movements may occur in the substructures of the tracker, thereby necessitating regular updating of the detector geometry to describe the position, orientation, and curvature of the tracker modules. The process of determining the new geometry parameters is called tracker alignment. Tracker alignment is performed numerous times throughout the data-taking period using reconstructed tracks from collisions and cosmic rays data and further fine-tuned once the data collection is completed. The strategies for and the performance of the tracker alignment during Run 2 (2016-2018) will be presented, emphasising the ultimate accuracy achieved with the legacy reprocessing. The data-driven techniques used to derive the alignment parameters and the methods used to validate the alignment performance will be reviewed. Finally, the preparations of CMS towards tracker alignment during Run 3 (2022-2025)–in particular, the very first alignment performed after the LHC Long Shutdown 2, with cosmic ray muons and collision data at $\sqrt{s}$ = 900 GeV–will be discussed.

1. The CMS Tracker

The innermost, all-silicon component of the CMS detector, the tracker, comprises two sub-detectors - the pixel detector and the strip detector [1]. The pixel detector, consisting of a barrel region (BPIX) and two forward endcaps (FPIX), is the closest in proximity to the beam interaction point. Until the end of the 2016 data-taking period, the BPIX of the so-called "Phase-0" pixel detector consisted of three layers and the FPIX endcaps each consisted of two disks, accounting for a total of 1440 modules. The upgraded "Phase-1" pixel detector, in operation since 2017, features an additional barrel layer and one more disk in each FPIX endcap with an increase in the number of modules to 1856 [2]. The strip detector, surrounding the pixels, contains 15 148 modules and consists of four sub-systems: the Tracker Inner Barrel and Disks (TIB and TID, respectively), the Tracker Outer Barrel (TOB), and the Tracker Endcaps (TEC). The arrangement of the modules in the tracker is such that they form several high-level structures, for example, two half barrels in the BPIX, four half-cylinders in the two FPIX regions, etc.

The purpose of the tracker is the precise determination of the trajectory of charged particles (tracks) from signals (hits), referred to as tracking. Robust tracking and detailed vertex reconstruction are instrumental in exploiting the range of physics accessible at the LHC.

Measuring the track momenta requires an accurate determination of the curvature of the tracks induced by the magnetic field, which in turn demands a well-aligned detector.

2. Alignment of the CMS Tracker

To achieve optimal track parameter resolutions, it is imperative to precisely measure each silicon sensor's position, orientation, and surface deformations. During the installation procedure, the module insertion constraints result in precision in the position of the tracker of $\mathcal{O}(0.1\,\mathrm{mm})$, which is larger than the design hit resolution of $\mathcal{O}(0.01\,\mathrm{mm})$ by one order of magnitude. Therefore, a further correction needs to be derived to push the alignment precision well below 0.01 mm. This correction is commonly referred to as tracker alignment, and the parameters of this correction as alignables.

CMS performs the alignment of the tracker using a track-based alignment method [3]. Every hit registered in the detector is assigned a measured hit position, and a set of tracks is formed from the combination of these hits. Each of these tracks is assigned a unique set of track parameters as well as alignment parameters. The track-based alignment method then follows a least-square approach to derive alignables $\mathbf{p}$ by minimising the following χ^2 function:

$$\chi^2(\mathbf{p},\mathbf{q}) = \sum_{j}^{\text{tracks}} \sum_{i}^{\text{hits}} \left(\frac{m_{ij} - f_{ij}\left(\mathbf{p},\mathbf{q}_j\right)}{\sigma_{ij}^{\mathrm{m}}} \right)^2 \tag{1}$$

where

- $\mathbf{q}$ represents the track parameters (e.g. parameters related to the track curvature and the deflection by multiple scattering),
- m represents the measured hit position and f the predicted hit position, and
- σ^{m} represents the uncertainty in measurement.

This minimisation can be carried out either globally using MillePede-II [3, 4], or locally, using HipPy [5, 6]. Subsequently, the tracks are refitted assuming the geometry defined by the updated set of alignables and the track momentum measurement is corrected. This is followed by validation of the new alignment conditions.

3. Tracker Alignment Strategy for Run 2

Every year during Run 2 (2016-2018), before the start of LHC collisions, at least the high-level structures of the tracker were aligned using the cosmic ray data available from the commissioning of the detector. This initial alignment allowed a preliminary coarse determination of the alignables, which can reveal substantial shifts (due to temperature and magnetic field changes, or the reinstallation of detector components during the detector shutdown, etc.) with regard to the initially assumed geometry. These alignment parameters were then used online during the data taking. The movements of the high-level structures of the pixel detector are continuously monitored by an automated alignment, which corrects the geometry if the alignment corrections exceed specific thresholds. Alongside, track-based alignments were also run offline regularly at different granularities of the tracker– from module-level to high-level structures. The automated alignment was then refined with regular updates from these offline computations. At the end of the data-taking period, the complete statistics of the dataset collected during the year were exploited to extract the alignment conditions and perform the "end-of-year (EOY) reconstruction". After the completion of Run 2, the ultimate accuracy of the alignment calibration was derived using the data collected over the three years and used for the final or legacy reprocessing of the data. Over these three years, about 700 000 parameters and 220 geometries, for various intervals of validity (IOVs), were calculated, covering the considerable variations in alignment circumstances over time.

4. Run 2 Results

The performance of different sets of alignment parameters derived during the data-taking, EOY reconstruction and legacy reprocessing [7] are illustrated in this section. The validation procedures used to compare the quality of these sets of alignables are also explained.

4.1. Tracking Performance

A good measure of the tracking performance is the distribution of median track-hit residuals (DMRs) per module. Each track is refitted, removing the hit under consideration to avoid any bias in the measurement. Ideally, the distribution should be narrow and centred at zero. The width of the DMR is influenced by random misalignment of the modules, while deviations of the mean from zero may be indicative of a systematic shift of the structure under scrutiny. The DMRs are calculated for all tracker substructures, as shown in Fig. 1. It is observed that after the dedicated alignment for the legacy reprocessing, the mean value μ of the DMRs is shifted closer to zero.

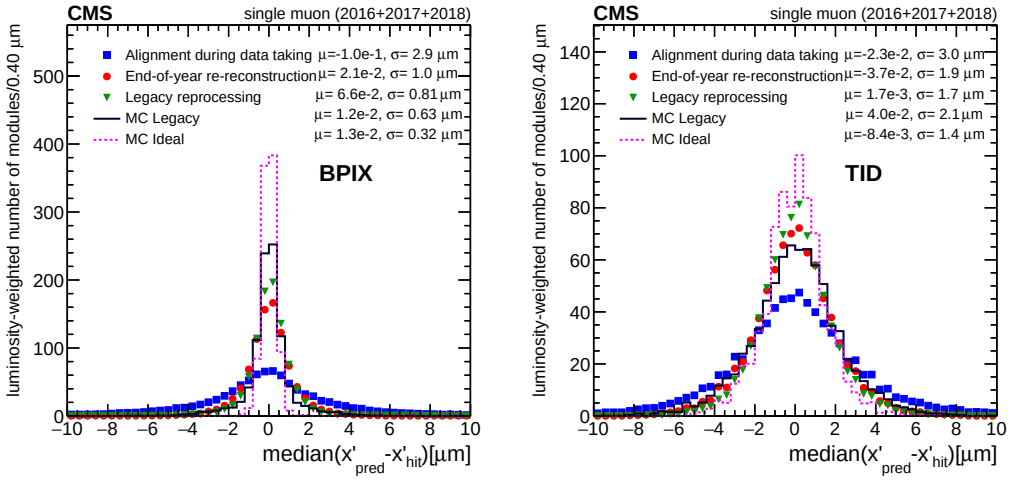

Figure 1. Distributions of median track-hit residuals per module in the local x (x') coordinate, for the BPIX (left) and TID (right). The distributions are averaged over all IOVs, where each IOV is weighted with the corresponding delivered integrated luminosity. The quoted means μ and standard deviations σ are the parameters of a Gaussian fit to the distributions.

4.2. Vertexing Performance

Studying the vertexing performance is essential to understand the effect of the alignment calibration on the reconstruction of physics objects. The unbiased track-vertex residuals, i.e. the distance between the tracks and the vertex reconstructed excluding the track under scrutiny, is examined to search for potential biases in the primary vertex reconstruction. As explained in the previous section, a deviation of the mean of the distribution from zero indicates a systematic misalignment. Fig. 2 shows the trends in the average unbiased track-vertex residuals in the transverse plane as a function of the delivered integrated luminosity, indicating an improved performance with legacy reprocessing.

4.3. Validation of Systematic Distortions

Validation is performed to check the influence of misalignments by projecting alignment performance onto a variable of interest. This variable has a known fixed value under perfect

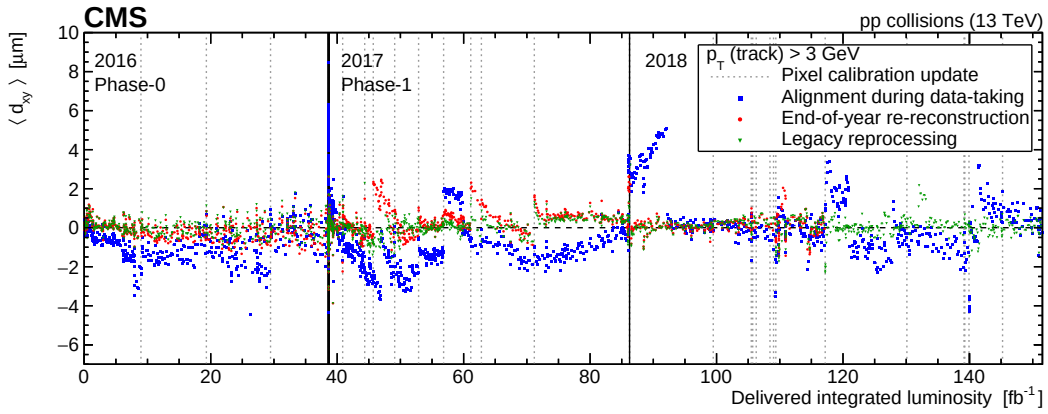

Figure 2. Average impact parameter trends in the transverse plane d_{xy} as a function of the delivered integrated luminosity.

conditions. In an ideally aligned tracker, the reconstructed $Z \to \mu\mu$ invariant mass, $m_{\mu\mu}$, should depend minimally on the direction in which the muons travel in the detector. Therefore, if the mean reconstructed mass differs from the expected m_Z value of 91.2 GeV, it indicates the presence of misalignment in the tracker. The uniformity in $m_{\mu\mu}$ that was achieved only after the legacy reprocessing can be seen in Fig. 3 (left). Cosmic ray muon tracks are another key ingredient used to control systematic distortions. The upper and lower portions of cosmic ray muon tracks that cross the tracker can be independently reconstructed, and the track parameters at the point of closest approach to the nominal beamline can be compared. Systematic differences between the track halves can indicate a misalignment, thus making this method a very powerful tool for validating track parameter resolutions. The mean of the difference in η between the two half-tracks refitted from the hits of a cosmic ray muon traversing the detector as a function of the impact parameter in the transverse plane is shown in Fig. 3 (right). Evidently, the strategy followed in the legacy alignment procedure has led to better performance.

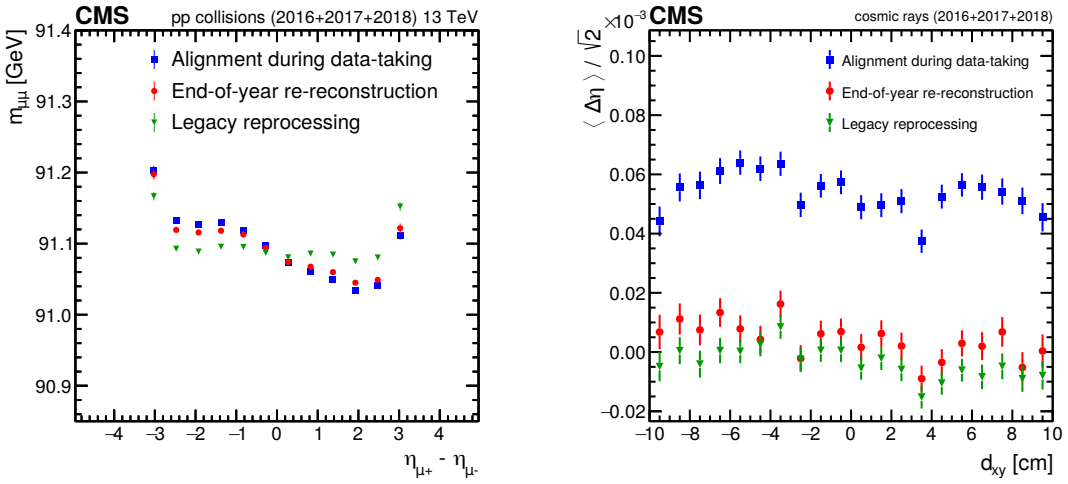

Figure 3. Reconstructed Z boson mass as a function of the difference in η between the positively and negatively charged muons, calculated from the full sample of dimuon events in Run 2 (left). The mean η difference between the two halves of cosmic muon tracks as a function of d_{xy}, scaled down by $\sqrt{2}$ to account for the two independent measurements (right).

5. LHC Run 3

5.1. Commissioning of the CMS Detector

The LHC underwent its second long shutdown (LS2) from 2018 to 2021. During this period, the BPIX of the CMS tracker was disassembled and its innermost layer was entirely replaced. Commissioning and calibrations are done, using cosmic ray muons, to ensure the CMS detector's proper working after such changes. Prior to the start of the LHC collisions in 2022, events were recorded during the *Cosmic RUns at ZEro Tesla* (CRUZET) before turning on the magnetic field from July to August 2021. Cosmic ray muon tracks were also recorded in the 3.8 T magnetic field provided by the CMS solenoid during the *Cosmic Runs At Four Tesla* (CRAFT), along with collisions at 900 GeV. These muon tracks are crucial for the alignment procedure as they can be employed to derive the first alignment corrections after the shutdown period. Also, they can be used to constrain several systematic distortions as described in section 4.3. The DMRs obtained after the alignment following CRUZET and CRAFT is shown in Fig. 4.

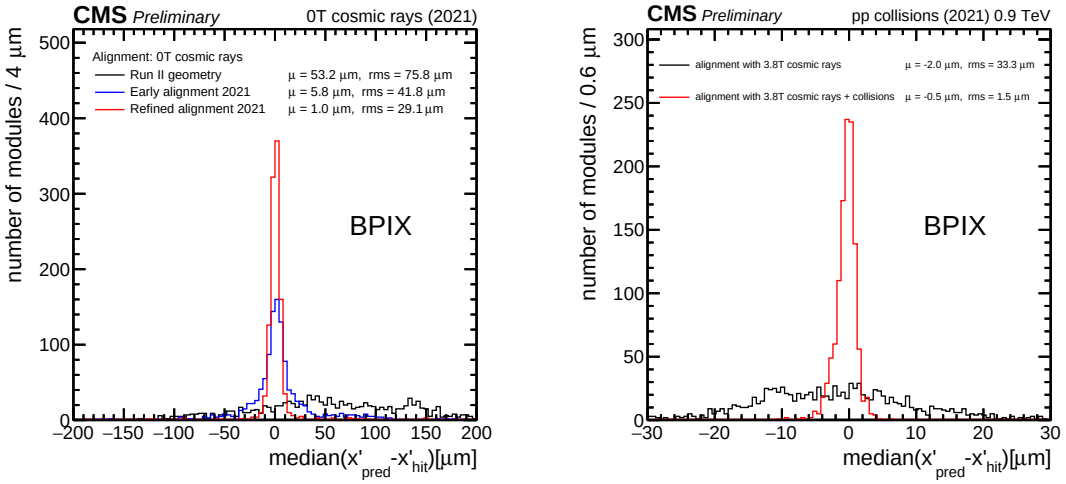

Figure 4. Distribution of median residuals in barrel pixel (BPIX) modules, along the local-x (x') direction obtained using CRUZET [8] (left) and CRAFT [9] (right) data.

5.2. Alignment Prospects

One of the significant goals during Run 3 would be to deploy finer granularity for the automated alignment. This run's larger irradiation doses will cause more substantial variation in the Lorentz drift of charge carriers in the tracker. Even though the alignment procedure is sensitive to Lorentz drift changes induced by accumulated radiation after $\approx 1\,\mathrm{fb}^{-1}$, the pixel local reconstruction calibration that corrects for this effect is performed only after $\approx 10\,\mathrm{fb}^{-1}$. If the alignment is done at fine enough granularity, inward and outward-pointing modules can move independently, absorbing the bias coming from Lorentz angle miscalibration. The alignments run offline will also be performed at a finer granularity to cope with the radiation effects.

6. Summary

The strategies and data-driven methods used to derive the alignment parameters for the CMS tracker during the LHC Run 2 were described. The performance of the tracker alignment after the legacy reprocessing was compared to the performances during data-taking and after EOY reconstruction. Systematic distortions arising in the detector geometry were studied using specific distributions. Finally, the commissioning and current status, as well as prospects of the alignment procedure for LHC Run 3, were discussed.

References

[1] The CMS Collaboration 2008 The CMS Experiment at the CERN LHC *JINST* **3** S08004 doi:10.1088/1748-0221/3/08/S08004
[2] The CMS Collaboration 2012 CMS Technical Design Report for the Pixel Detector Upgrade CERN-LHCC-2012-016 CMS-TDR-011 doi:10.2172/1151650
[3] Volker Blobel and Claus Kleinwort 2002 A New Method for the High-Precision Alignment of Track Detectors (*arXiv* hep-ex/0208021)
[4] The CMS Collaboration 2014 Alignment of the CMS tracker with LHC and cosmic ray data *JINST* **9** P06009 doi:10.1088/1748-0221/9/06/p06009
[5] Karimäki V, Lampen T and Schilling F P 2006 The HIP Algorithm for Track Based Alignment and its Application to the CMS Pixel Detector CMS-NOTE-2006-018 https://cds.cern.ch/record/926537
[6] Brown D, Gritsan A, Guo Z and Roberts D 2009 Local alignment of the BaBar Silicon Vertex Tracking detector *NIM-A* **603** 467–484 doi:10.1016/j.nima.2009.02.001 ISSN 0168-9002
[7] The CMS Collaboration 2021 Strategies and performance of the CMS silicon tracker alignment during LHC Run 2 (*arXiv* 2111.08757)
[8] The CMS Collaboration 2021 First Tracker Alignment results with 2021 cosmic ray data https://cds.cern.ch/record/2781754
[9] The CMS Collaboration 2021 148th LHCC Meeting - Open Session https://indico.cern.ch/event/1091297/contributions/4588574/attachments/2347421

Appendix C

Automated alignment calibration in CMS during Run 3

This section discusses the automated alignment workflow introduced in mid-2022 during Run 3 of the LHC. This workflow updates the tracker detector geometry by automatically processing the latest data and refining the alignment parameters within 48 hours. Improvements in alignment accuracy resulting from this continuous alignment are also highlighted. The content is presented as originally published in [229].

PROCEEDINGS OF SCIENCE

Automated alignment calibration in CMS during Run 3

Lakshmi Priya Nair[a,*] **for the CMS Collaboration**

[a] *Deutsches Elektronen-Synchrotron,*
Notkestraße 85, 22607 Hamburg, Germany

E-mail: lakshmi.priya.sreelatha.pramod@cern.ch

The LHC physics program requires a robust and efficient reconstruction of the trajectories of charged particles, as well as precise measurement of primary and secondary vertices and impact parameters. Several changes during data-taking, including the radiation damage introduced by the high particle flux, influence the particle position measurements. In order to exploit the physics potential of the Compact Muon Solenoid (CMS) detector, the calibration challenge consists of providing calibration constants with a fast turnaround, addressing the changes in the running conditions to ensure an efficient online event selection by the high-level trigger system and good quality first reconstruction of physics objects. To achieve this, several automated workflows have been developed. The update of the tracker detector geometry is performed by one such automated workflow, running on the most recent data and regularly updating the so-called alignment parameters within 48 hours. Thus, changes to the position and orientation of the tracker modules are periodically accounted for in the reconstruction. This contribution reviews the design and operational experience of the automated alignment calibration in place for the alignment of the CMS tracker detector, with an emphasis on the recent developments for the Run 3 data-taking period.

The European Physical Society Conference on High Energy Physics (EPS-HEP2023)
21-25 August 2023
Hamburg, Germany

*Speaker

PoS(EPS-HEP2023)595

1. Automation of calibrations in CMS

The Compact Muon Solenoid (CMS) is a general-purpose detector operated at the Large Hadron Collider (LHC), comprising different subdetectors – the pixel and strip tracker, the electromagnetic calorimeter (ECAL), the hadron calorimeter (HCAL) and the muon chambers – each performing a specific task. A detailed description of the CMS detector can be found in Ref. [1]. To make full use of the physics potential of the CMS experiment, it is paramount to calibrate and align the different subdetectors from time to time.

One of the critical assets for detector operations is having the most accurate calibrations available with a fast turnaround. This enables the delivery of good-quality data that can be used for physics analyses within 48-72 hours of their acquisition. Furthermore, the need to perform additional data reprocessing is reduced by ensuring improved calibration for initial physics object reconstruction. An elaborate framework called the Prompt Calibration Loop (PCL) was thus developed to fulfil this need for fast calibrations [2]. The PCL includes several low-latency alignment and calibration workflows that are run automatically at the Tier-0 processing farm at CERN for each acquired run. The strategy for running these is based on the 48-hour delay between the data acquisition and the reconstruction of the bulk of this data for physics analyses, called prompt reconstruction. Figure 1 shows a schematic illustrating the steps used by the different PCL workflows to compute alignment and calibration constants. Within 1-2 hours of the data collection, a limited selection of these data, called the express stream, is first reconstructed to provide rapid feedback on detector status and physics performance and serves as input to calibration workflows. The express data is further skimmed by customising the event content for each Alignment and Calibration (AlCa) workflow, thus producing AlCaRECO datasets. The calibration algorithms are then executed in parallel, using the AlCaRECO datasets as inputs to produce intermediate calibration products (histograms or calibration constants computed by a single job). The last step, called AlCa Harvesting, involves the aggregation of all the intermediate products for a given run into a set of conditions (database payloads) for future consumption by the reconstruction jobs. Initially stored as SQLite files, the payloads are then transferred to the Oracle-based CMS AlCa database [3]. In addition, histograms are also produced and loaded into an instance of the Data Quality Monitoring (DQM) [4] framework to allow a review of the performance of calibration algorithms.

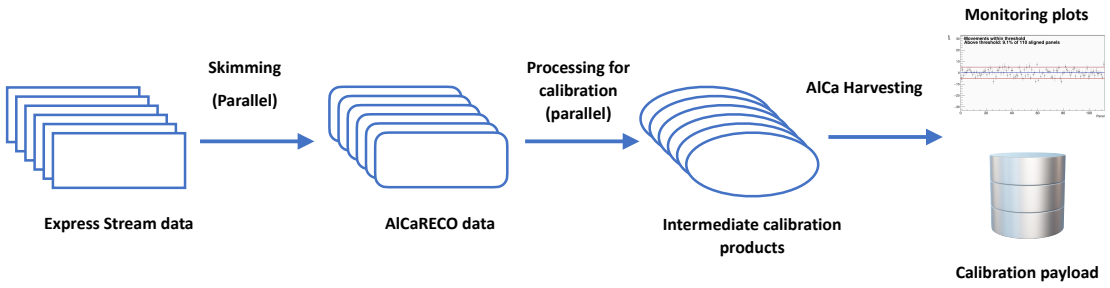

Figure 1: Schematic view of the working of the automated alignment and calibration workflows to compute conditions for one run.

2. Alignment of the CMS tracker

The CMS tracker comprises two sub-detectors - the silicon pixel and silicon strip detectors. The pixel detector, consisting of a barrel region (BPIX) and two forward endcaps (FPIX), with a total of 1 856 modules [5], is the closest to the beam interaction point. The strip detector surrounding the pixels contains 15 148 modules. The CMS tracker is designed for precise tracking, i.e. determining the trajectories of charged particles (tracks) from signals (hits) and accurate vertex reconstruction. These, in turn, demand a well-aligned detector where each sensor's position, orientation, and surface deformation are precisely known. The tracker has a design hit resolution of $O(10\,\mu\text{m})$. Therefore, corrections need to be derived to push the alignment precision well below this resolution. The track-based alignment method used to do this follows a least squares approach to derive alignables $\mathbf{p}$ by minimising the following χ^2 function:

$$\chi^2(\mathbf{p},\mathbf{q}) = \sum_{j}^{\text{tracks}} \sum_{i}^{\text{measurements}} \left(\frac{m_{ij} - f_{ij}(\mathbf{p},\mathbf{q})}{\sigma_{ij}} \right) \tag{1}$$

where $\mathbf{q}$ represents the track parameters and the track-hit residual $(m_{ij} - f_{ij})$ is obtained by subtracting the hit prediction (f_{ij}) from the measured hit position (m_{ij}) [6]. σ_{ij} corresponds to the uncertainty in m_{ij}. This minimisation can be done either globally using the MillePede-II [7] or locally using the HipPy algorithm [8].

3. Automated tracker alignment

During data-taking, the different components of both the pixel and strip tracker may shift because of changes in the magnetic field or the temperature. Being closest to the interaction point, the pixel detector is subjected to much higher levels of radiation than the strip detector, the effects of which also have to be taken into account when deriving alignment parameters. As changes occur frequently in the pixels, an automated alignment procedure is needed to correct them. Thus, the automated alignment workflow, now called the low-granularity PCL (LG-PCL), was first implemented as part of the PCL during Run 2 of LHC. It works as explained in Section 1 and accounts for the movement of the high-level structures in the pixel detector, i.e. two half-barrels and two half-cylinders. It derives corrections using the MillePede-II alignment algorithm based on 36 degrees of freedom - the positions (x, y, z) as well as the rotations $(\theta_x, \theta_y, \theta_z)$ for each of the structures.

The LG-PCL, active throughout Run 2, showed stable performance during this period and helped correct large movements caused by magnet cycles, which is explained in more detail in Ref. [8]. However, it has a few shortcomings. A quantity sensitive to the radiation dose which plays a role in alignment is the Lorentz drift, which depends on the electric field and the mobility of the charge carriers in the silicon sensors, among other factors. Because of the high radiation dose, the mobility of the charge carriers changes quickly in the pixel detector. If not accounted for properly, this will cause systematic biases. This is corrected using a dedicated local calibration method, and the residual effects need to be corrected in the alignment procedure. The effectiveness of this can be monitored by producing the distribution of medians of track-hit residuals (DMR) separately for the inward- and outward-facing modules (which have different directions of electric fields and thus

different directions of the Lorentz drift) and calculating the differences between the means of the DMRs, $\Delta\mu$. Ideally, the value of $\Delta\mu$ must be zero, and a different value hints at residual biases due to the accumulated effects from radiation in the sensors. Figure 2 shows the $\Delta\mu$ distributions for the BPIX for the years 2016-2018 as a function of the delivered integrated luminosity. Though successful on several fronts, the LG-PCL (blue curve) cannot fully mitigate the radiation effects, as the granularity of the alignment parameters used is too coarse to account for these changes. The bias could only be treated during the legacy reprocessing of Run 2 data by using additional parameters.

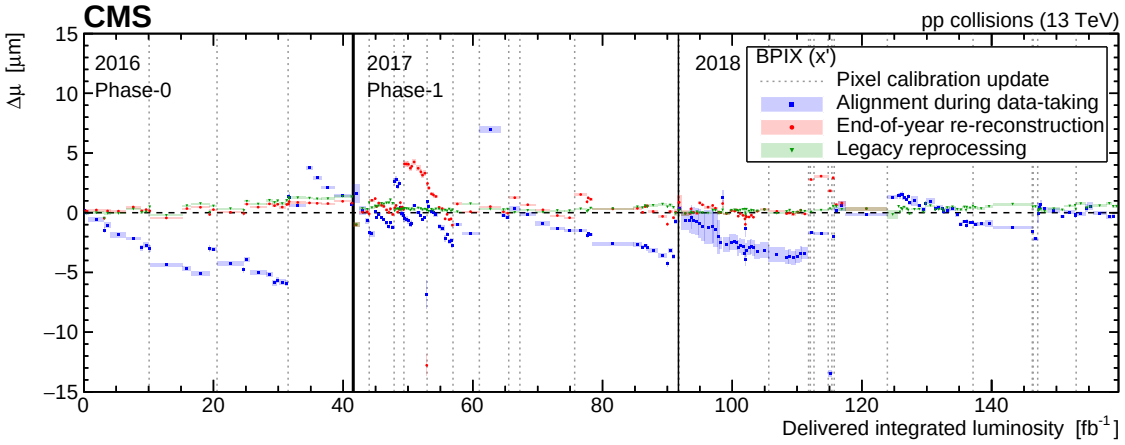

Figure 2: The difference, $\Delta\mu$, between the mean values of the DMRs obtained separately for the modules with the electric field pointing radially inwards or outwards for the local x (x') coordinate in the BPIX detector for the years 2016-2018, as a function of the delivered integrated luminosity [8].

The above is one of the main motivations behind developing the new automated alignment workflow, referred to as the high-granularity PCL (HG-PCL), for Run 3. The HG-PCL also uses the MillPede-II algorithm and is implemented in the central PCL workflow. It aligns the individual ladders and panels of the pixel detector, thus operating at a finer granularity than the LG-PCL and aligning over 5 000 parameters. After extensive testing and validation, the HG-PCL was deployed for data-taking in September 2022 and is successfully providing alignment conditions since then. The HG-PCL reduces the systematic bias induced by radiation damage as the rapidly changing shift from the local reconstruction can be absorbed in the position of the ladders and panels. This can be clearly observed in Figure 3, which shows $\Delta\mu$ for the innermost barrel layer, exposed to the harshest radiation environment, as a function of the delivered integrated luminosity for the year 2022. The blue curve shows the performance of the alignment conditions produced by the HG-PCL, which were also used for the end-of-year reconstruction due to their excellent performance, which allowed to save a lot of time and resources.

4. Summary

A description of the design and operation of the automated alignment calibration of the CMS tracker to cope with the rapid changes in data-taking conditions was given. Developments made in the workflows aiming to improve the alignment conditions used already at the level of prompt reconstruction during Run 3 were discussed. The promising results achieved thus far were presented.

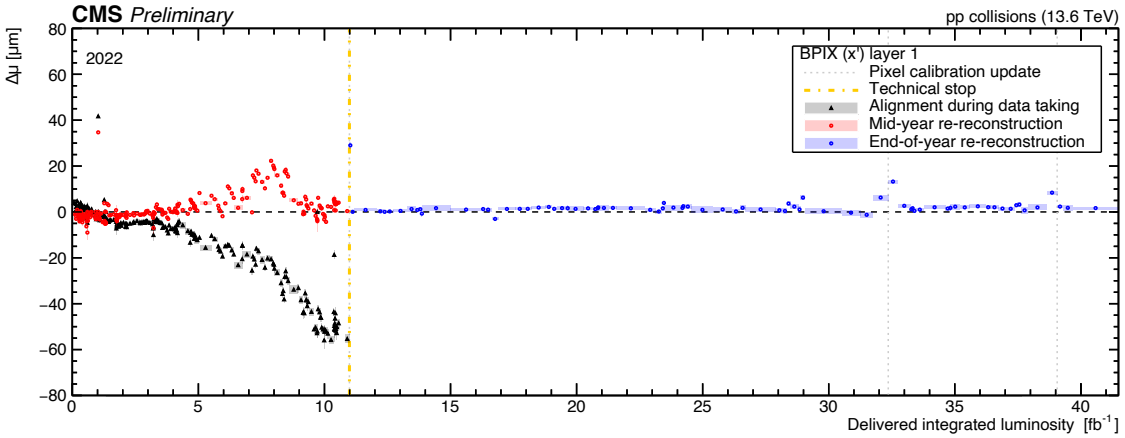

Figure 3: The difference, $\Delta\mu$, between the mean values of the DMRs obtained separately for the modules with the electric field pointing radially inwards or outwards for the local x (x') coordinate in BPIX layer 1 for the year 2022, as a function of the delivered integrated luminosity [9]. The blue curve shows the performance of the alignment constants derived by the HG-PCL.

References

[1] CMS Collaboration, *The CMS Experiment at the CERN LHC*, *JINST* **3** (2008) S08004.

[2] G. Cerminara and B. van Besien, *Automated workflows for critical time-dependent calibrations at the CMS experiment*, *J. Phys. Conf. Ser.* **664** (2015) 072009.

[3] CMS Collaboration, *CMS experience with online and offline databases*, *J. Phys. Conf. Ser.* **396** (2012) 052059.

[4] L. Tuura, A. Meyer, I. Segoni and G.D. Ricca, *CMS data quality monitoring: Systems and experiences*, *J. Phys. Conf. Ser.* **219** (2010) 072020.

[5] CMS Collaboration, *CMS Technical Design Report for the Pixel Detector Upgrade*, CMS Technical Design Report CMS-TDR-11 (2012).

[6] CMS Collaboration, *Alignment of the CMS tracker with LHC and cosmic ray data*, *JINST* **9** (2014) P06009 [1403.2286].

[7] V. Blobel and C. Kleinwort, *A New method for the high precision alignment of track detectors*, in *Conference on Advanced Statistical Techniques in Particle Physics*, 6, 2002 [hep-ex/0208021].

[8] CMS Collaboration, *Strategies and performance of the CMS silicon tracker alignment during LHC Run 2*, *Nucl. Instrum. Meth. A* **1037** (2022) 166795 [2111.08757].

[9] CMS Collaboration, *Tracker alignment performance in 2022 (addendum)*, CMS Detector Performance Summary CMS-DP-2022-070 (2022).

Appendix D

Double-muon Trigger Efficiencies in 2016 and 2017

In this section, the efficiencies calculated for the p_{T} legs of the 2016 and 2017 trigger, as explained in Sec. 5.6.6, are presented. Figures. D.1 and D.2 show the efficiencies corresponding to the higher-p_{T} and lower-p_{T} leg of the `HLT_Mu17_Mu8_SameSign` trigger. Figures D.3 and D.4 show the efficiencies of the higher-p_{T} and lower-p_{T} leg of the alternate `HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass8` trigger.

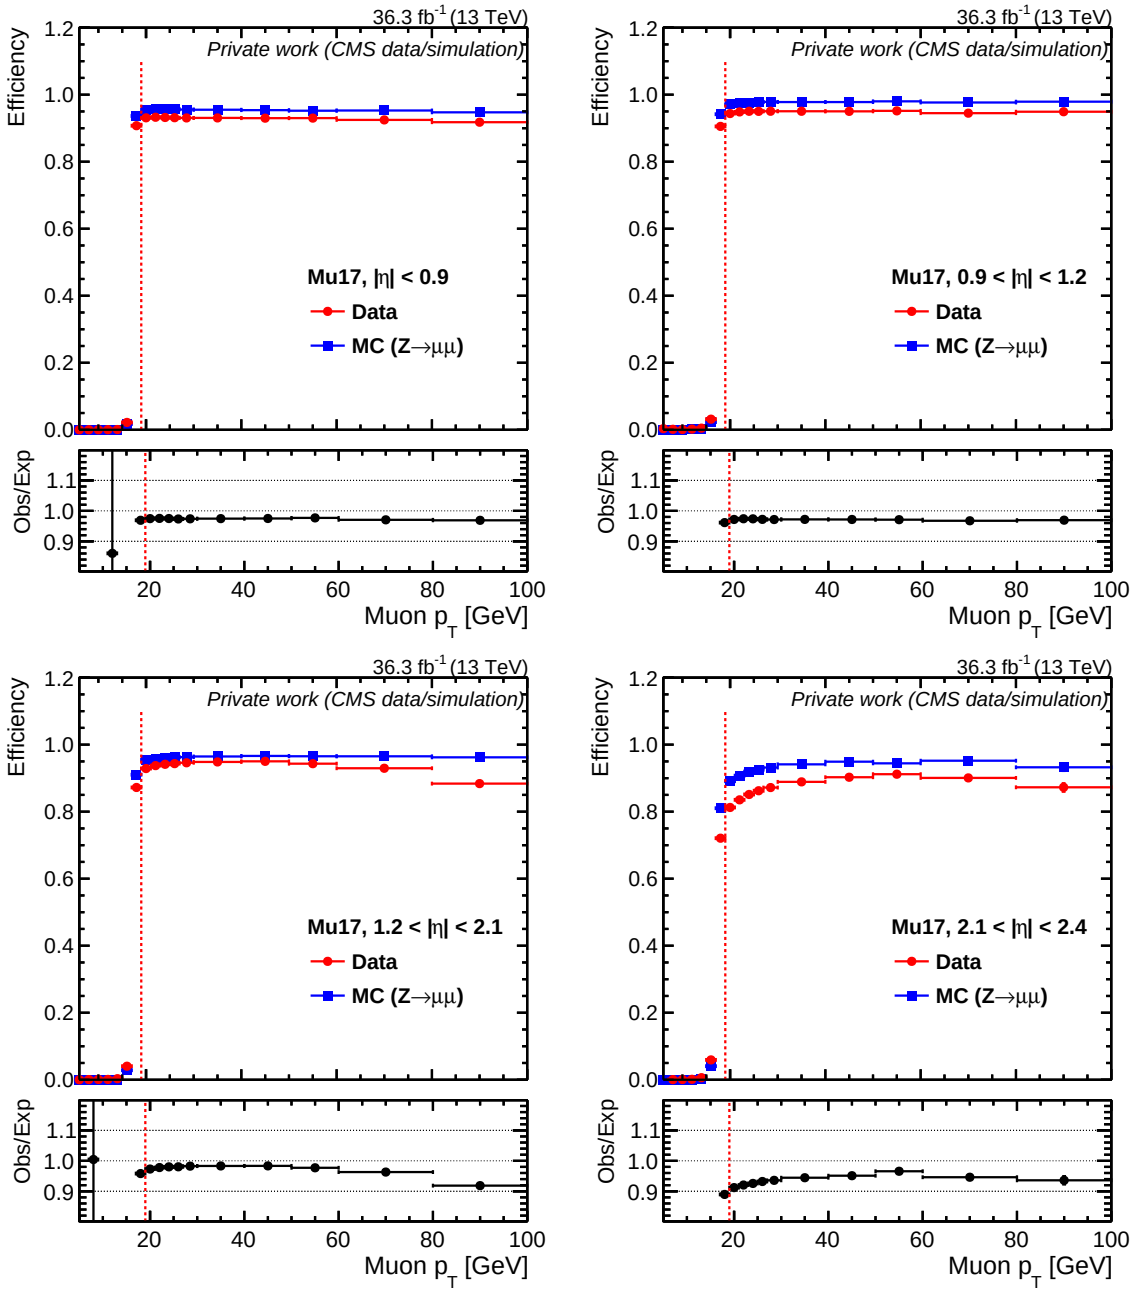

Figure D.1: Efficiency of the `Mu17` leg of the trigger `HLT_Mu17_Mu8_SameSign_DZ` in the 2016 data sample as a function of the muon transverse momentum for four different eta bins: $|\eta| < 0.9$ (top left), $0.9 < |\eta| < 1.2$ (top right), $1.2 < |\eta| < 2.1$ (bottom left), and $|\eta| > 2.1$ (bottom right). The circles present measurements in data, and the squares are measurements in MC.

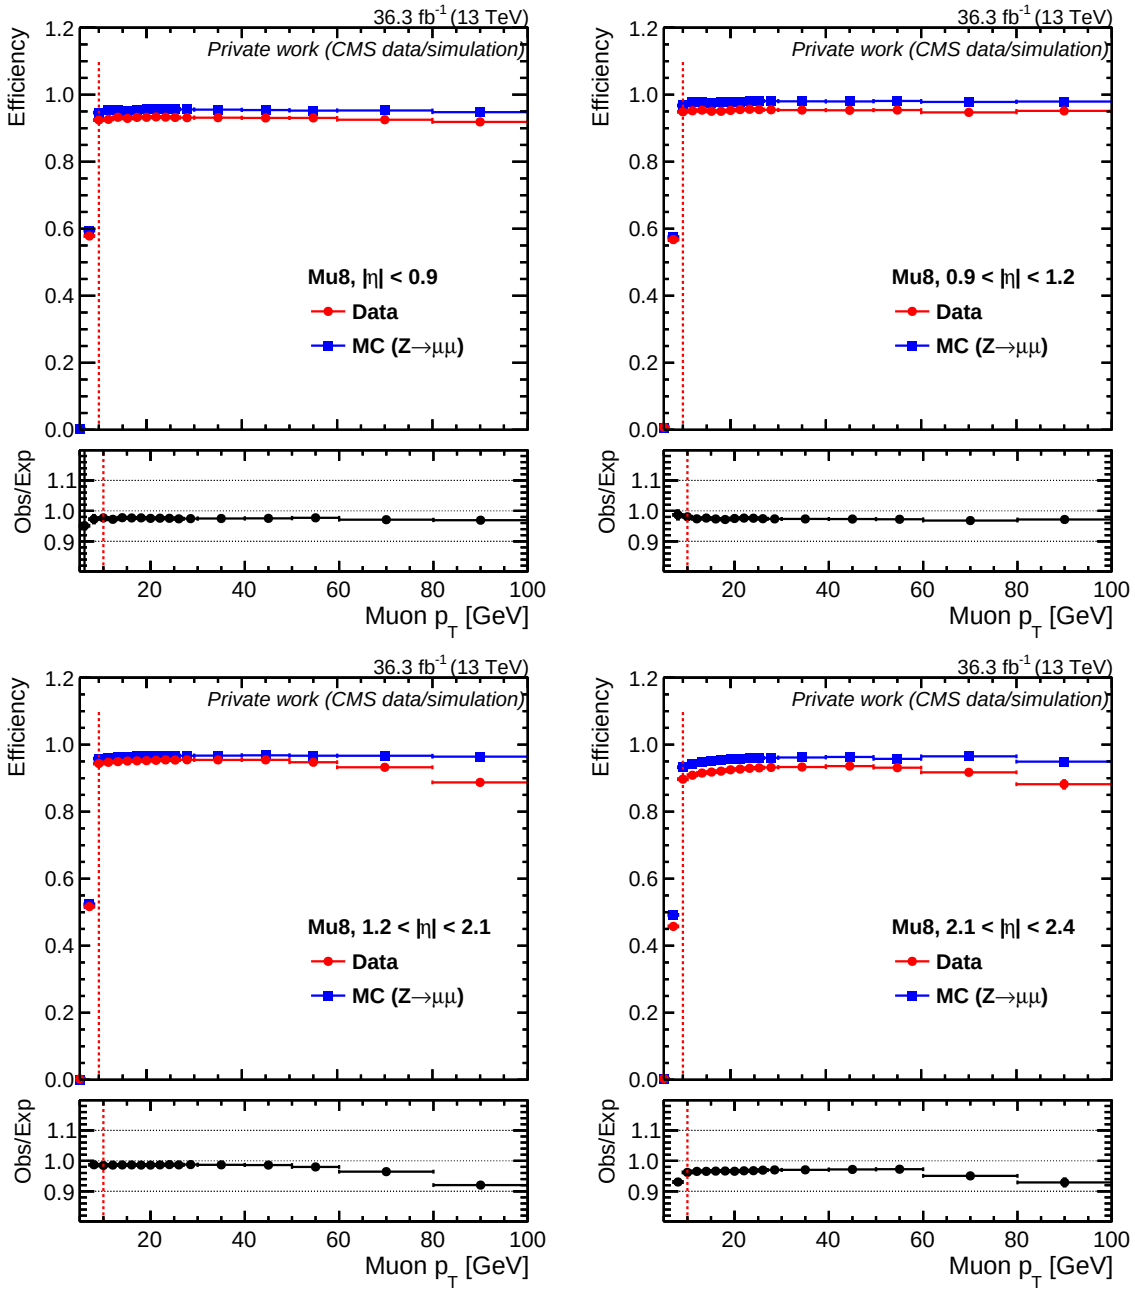

Figure D.2: Efficiency of the `Mu8` leg of the trigger `HLT_Mu17_Mu8_SameSign_DZ` in the 2016 data sample as a function of the muon transverse momentum for four different eta bins: $|\eta| < 0.9$ (top left), $0.9 < |\eta| < 1.2$ (top right), $1.2 < |\eta| < 2.1$ (bottom left), and $|\eta| > 2.1$ (bottom right). The circles present measurements in data, and the squares are measurements in MC.

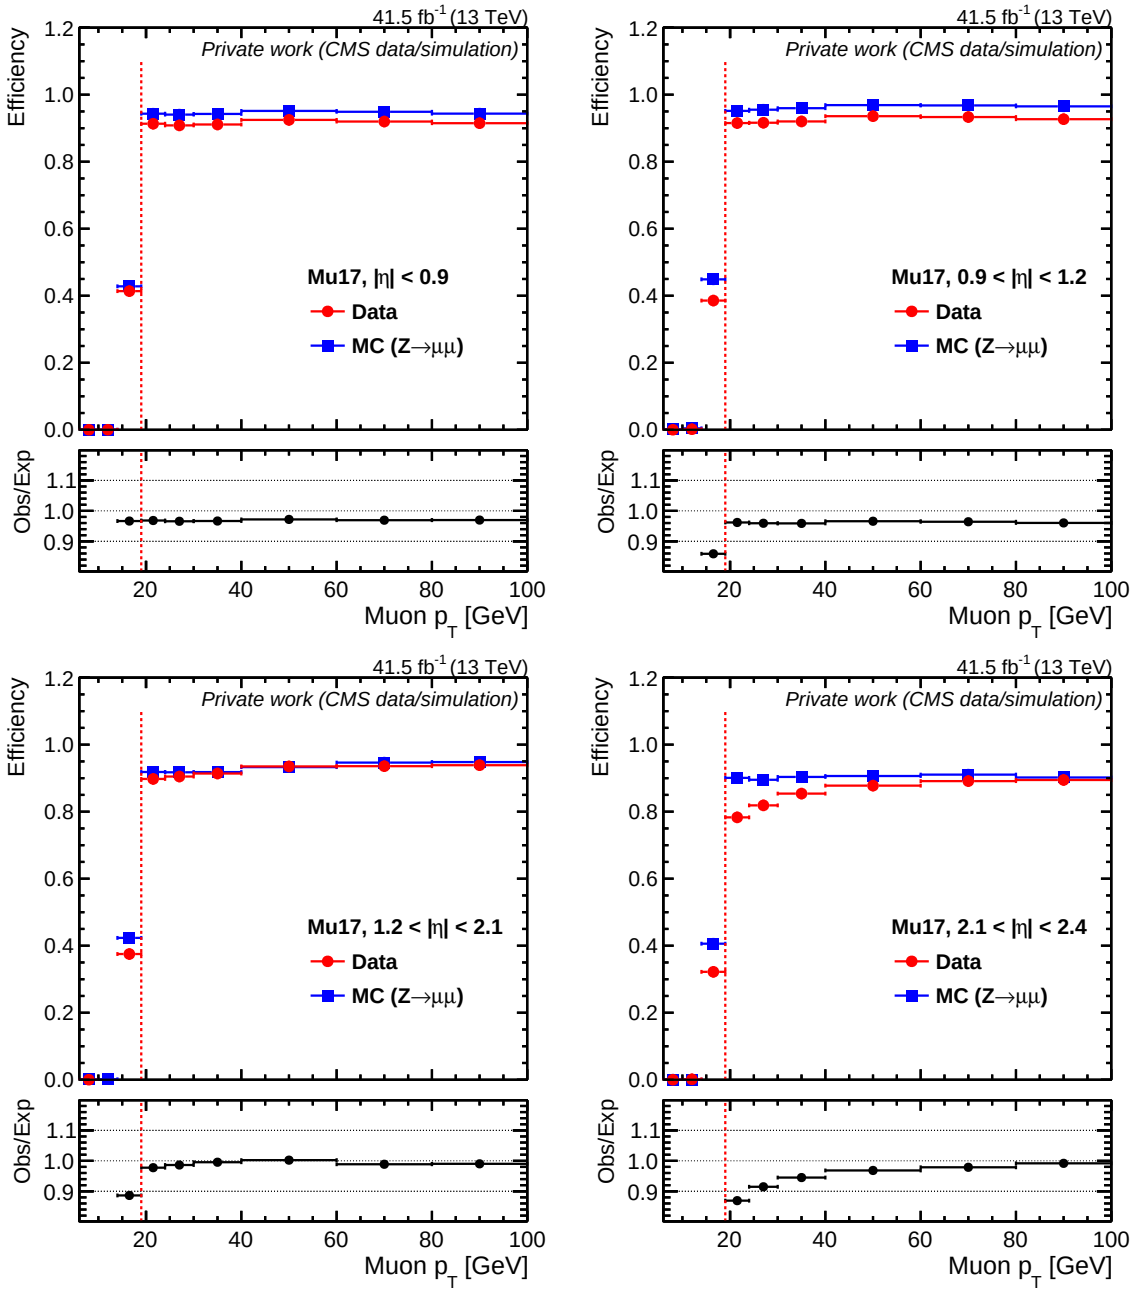

Figure D.3: Efficiency of the `Mu17` leg of the trigger `HLT_Mu17_TrkVVL_Mu8_TrkVVL_DZ_-Mass8` in the 2017 data sample as a function of the muon transverse momentum for four different eta bins: $|\eta| < 0.9$ (top left), $0.9 < |\eta| < 1.2$ (top right), $1.2 < |\eta| < 2.1$ (bottom left), and $|\eta| > 2.1$ (bottom right). The circles present measurements in data, and the squares are measurements in MC.

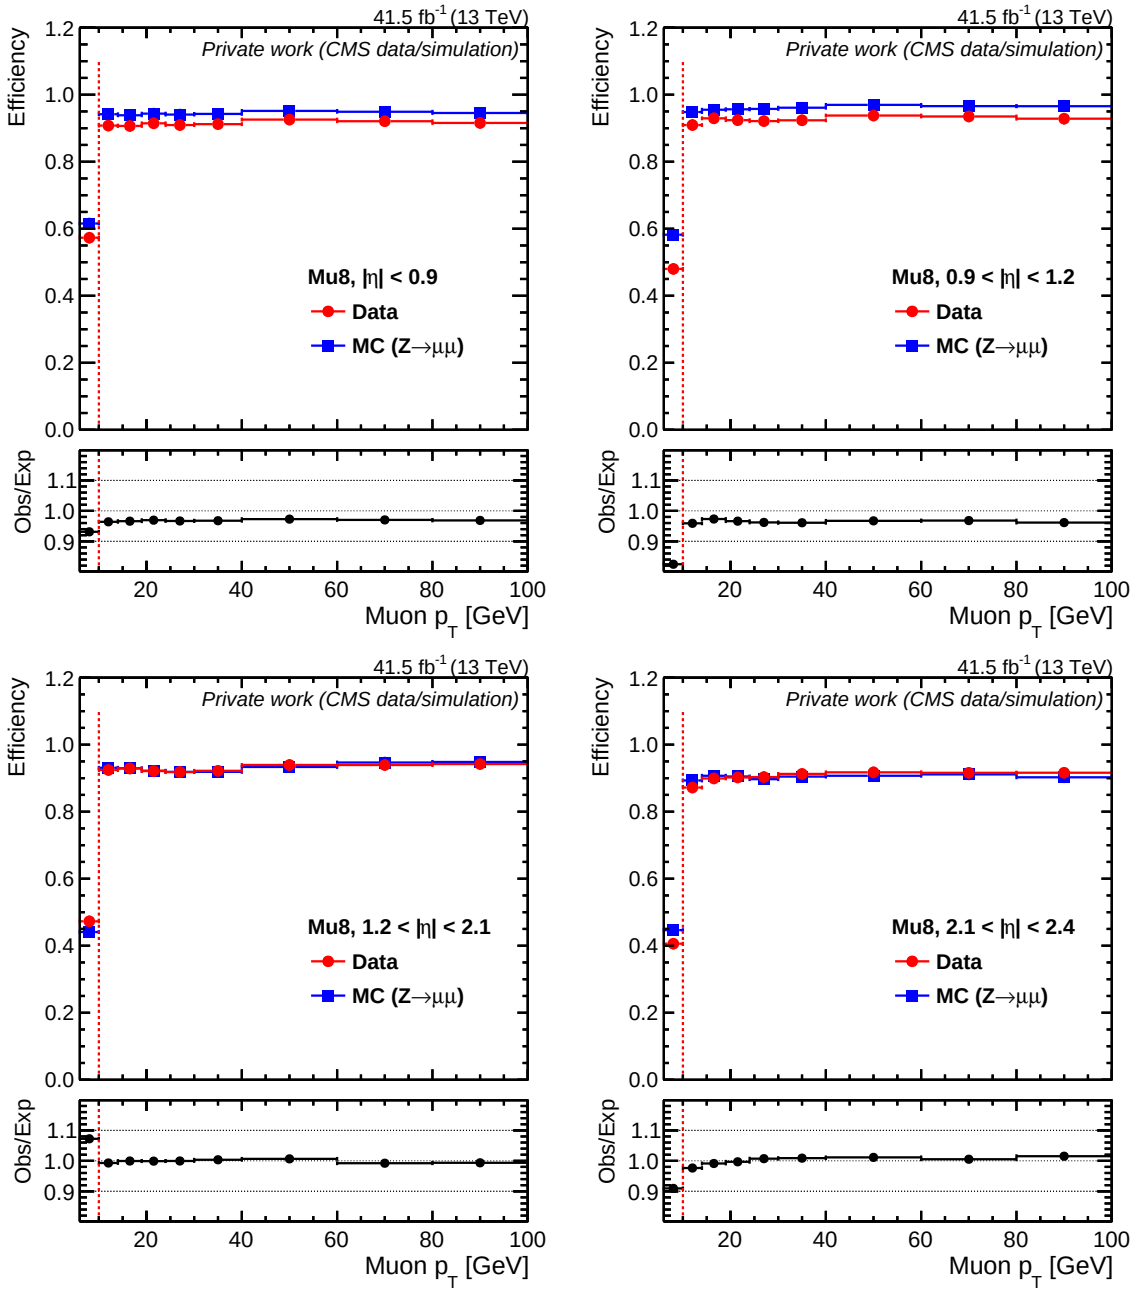

Figure D.4: Efficiency of the `Mu8` leg of the trigger `HLT_Mu17_TrkVVL_Mu8_TrkVVL_DZ_Mass8` in the 2017 data sample as a function of the muon transverse momentum for four different eta bins: $|\eta| < 0.9$ (top left), $0.9 < |\eta| < 1.2$ (top right), $1.2 < |\eta| < 2.1$ (bottom left), and $|\eta| > 2.1$ (bottom right). The circles present measurements in data, and the squares are measurements in MC.

Appendix E

Higgs-p_{T} Reweighting: Additional Plots

In Section 5.6.8, the reweighting applied to simulated signal samples to reflect higher-order predictions for the p_{T} distribution of the H boson was detailed. This section also showcased the NNLO p_{T} distribution of the Higgs boson for the gluon-gluon fusion production mode, which was obtained using the HQT (v.2.0) program. Furthermore, the k-factors were derived by dividing this NNLO distribution by the one from standard PYTHIA (v.8.212) samples. The plots for the other production modes—VBF, VH, and $\mathrm{t\bar{t}H}$—are presented here in Appendix E.

The higher-order predictions for the VBF Higgs production mode are obtained at NLO using the POWHEGBOX (v.2.0) generator. The corresponding distributions and k-factors are illustrated in Fig. E.1. For the VH production mode, predictions for the ZH, $\mathrm{W^{-}H}$, and $\mathrm{W^{+}H}$ modes are also made at NLO using POWHEGBOX. The resulting k-factors are displayed in Figs. E.2, E.3, and E.4, respectively. Lastly, for the $\mathrm{t\bar{t}H}$ mode, NLO k-factors are calculated using MADGRAPH5_AMC@NLO (v.2.6.5). These are shown in Fig. E.5.

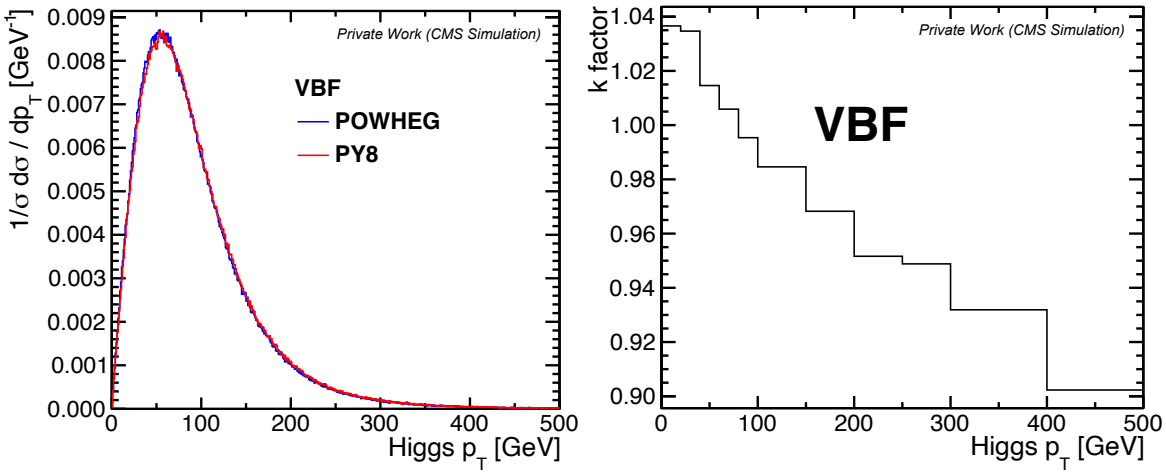

Figure E.1: The distribution of the Higgs p_T as predicted at LO by PYTHIA (v.8.212) and at NLO by POWHEGBOX (v.2.0) in the vector boson fusion process (left) and corresponding k-factor (right).

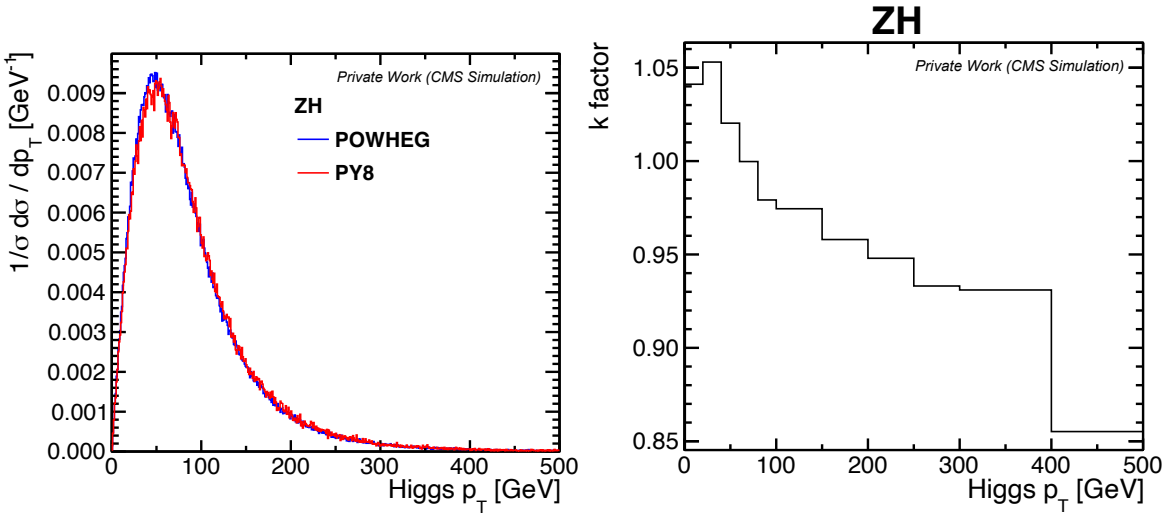

Figure E.2: The distribution of the Higgs p_T as predicted at LO by PYTHIA (v.8.212) and at NLO by POWHEGBOX (v.2.0) in the Z-associated production process (left) and corresponding k-factor (right).

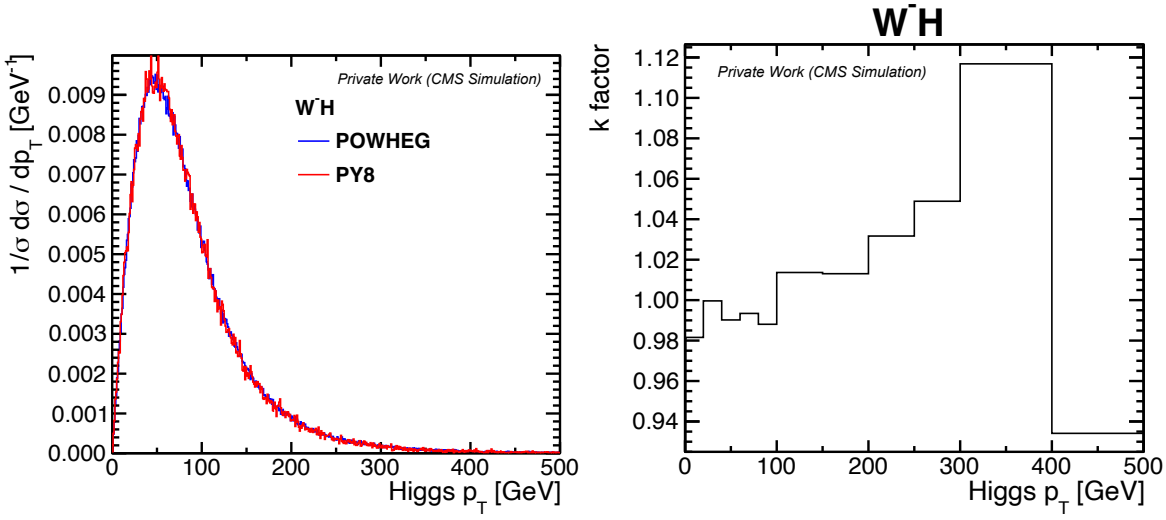

Figure E.3: The distribution of the Higgs p_{T} as predicted at LO by PYTHIA (v.8.212) and at NLO by POWHEGBOX (v.2.0) in the W^{-}-associated production process (left) and corresponding k-factor (right).

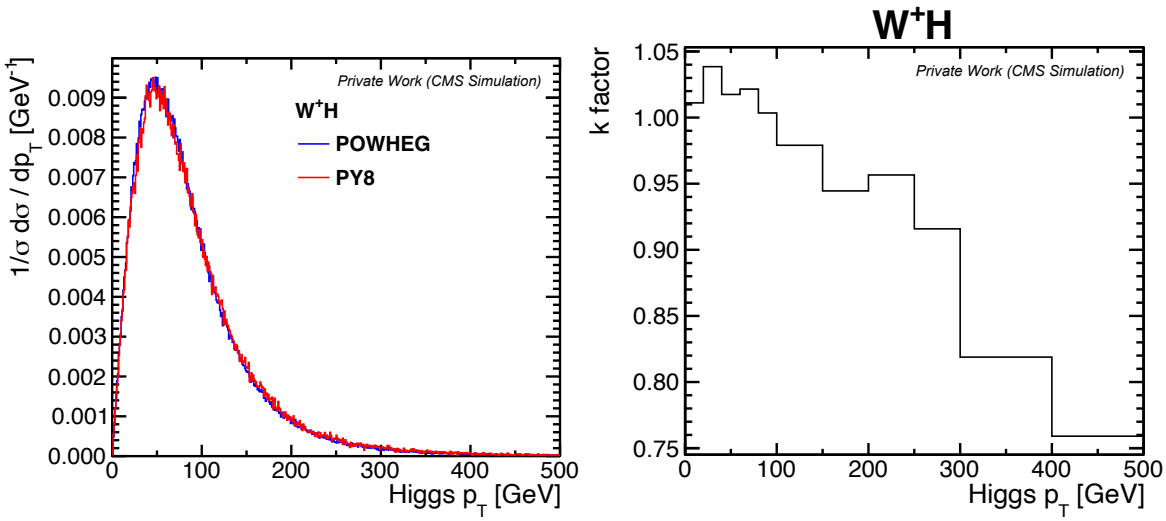

Figure E.4: The distribution of the Higgs p_{T} as predicted at LO by PYTHIA (v.8.212) and at NLO by POWHEGBOX (v.2.0) in the W^{+}-associated production process (left) and corresponding k-factor (right).

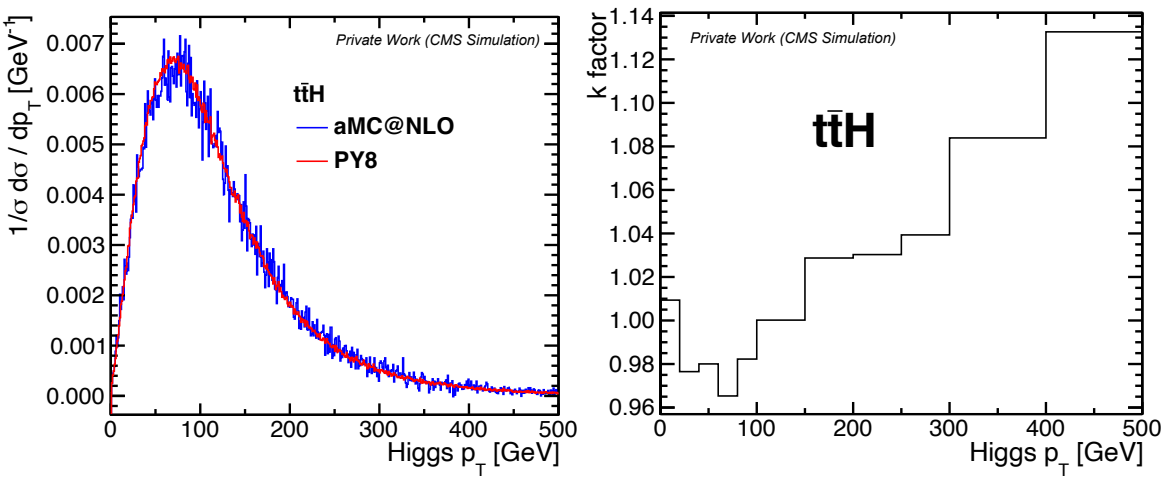

Figure E.5: The distribution of the Higgs p_T as predicted at LO by PYTHIA (v.8.212) and at NLO by MADGRAPH5_aMC@NLO (v.2.6.5) in the top-pair associated production process (left) and corresponding k-factor (right).

Appendix F

Measurement of Track Isolation and Identification Efficiency

Contents

F.1 Introduction **201**

F.2 Event Selection **202**

F.3 Extraction of Scale Factors **203**

F.4 Dependence of Scale Factors on $\Delta R_{\text{iso,trk}}$ **204**

F.1 Introduction

The muon-track isolation criterion requires the presence of only one track in the isolation cone around the muon momentum. As mentioned in Sec. 5.6.7, this isolation can be affected by charged particles originating from pileup interactions and underlying events. Conversion of photons that are radiated off muons or stem from the $\pi^0 \to \gamma\gamma$ decay affects isolation to a lesser extent. Differences in the isolation criteria can, in turn, also affect the identification of the one-prong tau candidates. As the pileup and underlying events modeled in simulated samples do not perfectly match those in the data, there would be differences between the track isolation and identification efficiencies obtained for the simulation and experimental data.

Hence, scale factors are derived for the combined efficiency of the track isolation and one-prong tau decay identification to correct the performance in simulated samples. The measurement is performed individually for the 2016, 2017, and 2018 data taking periods and is carried out using events containing genuine $\tau_\mu\tau_{\text{1-prong}}$ pairs. The largest contribution to these events comes from the $Z \to \tau_\mu\tau_{\text{1-prong}}$ events. However, other processes, such as the $t\bar{t}$, single-top and diboson productions, can also produce genuine $\tau_\mu\tau_{\text{1-prong}}$ final states.

F.2 Event Selection

Events of interest are selected with the `HLT_IsoMu24` trigger for 2016 and the `HLT_IsoMu27` trigger for 2017 and 2018. The events must contain a muon with $|\eta| < 2.4$ and $p_{\mathrm{T}} > 28\,\mathrm{GeV}$ if `HLT_IsoMu27` is used and $p_{\mathrm{T}} > 25\,\mathrm{GeV}$ if `HLT_IsoMu24` is used. The muon should pass the *medium* ID. The relative $\Delta\beta-$corrected particle-flow isolation of the muon is required to be $I_{\mu} < 0.15$.

Events are further selected if they contain an oppositely-charged track passing the one-prong tau identification criteria employed in the main analysis, i.e., the track must have a charge opposite to that of the muon, $p_{\mathrm{T}} > 2.5\,\mathrm{GeV}$ and $|\eta| < 2.4$. The polar angle between the muon and the one-prong tau candidate must fulfill $\Delta\phi(\mu, \mathrm{trk}) > 2.0$ rad. The isolation requirement for the track is identical to that used in the main analysis, except that the isolation cone is centered on the track momentum vector. Hence, the track is required to contain no other tracks with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$, $|\eta| < 2.4$, $|d_0| < 0.2$ cm, and $|d_z| < 0.3$ cm in an isolation cone of radius 0.5 around its momentum. If more than one track satisfies these selection criteria, the track that yields the maximal $\Delta\phi(\mu, \mathrm{trk})$ is selected.

The main background in this case includes events with genuine prompt muons and jets faking 1-prong taus, coming from all considered processes. Dimuon events coming from $\mathrm{Z} \to \mu\mu$ and QCD multijets, where both muon and one-prong tau candidates are produced within hadronic jets, also form considerable backgrounds. To reduce the contributions from $\mathrm{Z} \to \mu\mu$ decays, events with a second loosely identified muon are vetoed. To suppress top-pair, single-top, and diboson backgrounds, events are required to have zero jets. To suppress the W+jets background, the transverse mass of the muon and the missing transverse momentum are requested to satisfy $m_{\mathrm{T}} < 30\,\mathrm{GeV}$.

The above-mentioned selections define the signal region (SR) for this measurement.

The selected events from the SR are then classified into four regions, depending on the transverse momentum of the track:

- $5 < p_{\mathrm{T}} < 10\,\mathrm{GeV}$,
- $10 < p_{\mathrm{T}} < 15\,\mathrm{GeV}$,
- $15 < p_{\mathrm{T}} < 20\,\mathrm{GeV}$,
- $p_{\mathrm{T}} > 20\,\mathrm{GeV}$

A category for tracks in the p_{T} range 2.5–5 GeV is not considered, as in this range, the selected $\tau_{\mu}\tau_{\text{1-prong}}$ sample is entirely dominated by the backgrounds, and no reliable measurement is possible. Hence, the measurements are started at $p_{\mathrm{T}} = 5\,\mathrm{GeV}$ and extrapolated down to lower values, as explained in the next section.

F.3 Extraction of Scale Factors

The scale factor for the track isolation and identification efficiency is extracted by a simultaneous maximum likelihood fit of the muon-track invariant mass, $m_{\mu,\mathrm{trk}}$, in the four track-p_{T} regions of the selected $\tau_\mu\tau_{\text{1-prong}}$ sample. The prefit distributions of $m_{\mu,\mathrm{trk}}$ across various track-p_{T} ranges in the selected $\tau_\mu\tau_{\text{1-prong}}$, for example, from 2018 are shown in Fig. F.1.

Here, the shapes of $\mathrm{Z} \to \tau_\mu\tau_{\text{1-prong}}$, top-pair, single-top, single-W, diboson and $\mathrm{Z} \to \mu\mu$ templates are derived from simulation. The shape of the QCD multijet background is modeled using a sideband region, where the muon and the track have the same sign, by subtracting the other backgrounds from the data. Another control region, depleted in $\mathrm{Z} \to \tau_\mu\tau_{\text{1-prong}}$ events, containing events with a high transverse mass of muon and missing transverse energy ($m_{\mathrm{T}} > 50\,\mathrm{GeV}$), is used to constrain backgrounds with jets faking 1-prong tau. This region mainly consists of W+jets events. A third control region with $\mathrm{Z} \to \mu\mu$ events, orthogonal to the signal region, is used to constrain the normalization of Drell-Yan MC samples. In this region, a fit is performed in a single bin of the dimuon mass between 70 and 112 GeV.

The normalization of the QCD background is determined through an extrapolation factor from the same sign to the opposite sign region. The extrapolation factor is measured in the region of the anti-isolated muon individually for each track p_{T} bin and $\mathrm{m_T}$ region.

The likelihood fit is performed with four freely floating rate parameters:

1. The normalization of the simulated processes, yielding 1-prong taus in the final state, is scaled with a rate parameter, which is interpreted as the one-prong tau identification scale factor. This rate parameter is of primary interest.

2. The normalization of the simulated processes where jets are misidentified as 1-prong tau decays are scaled by a freely floating rate parameter. This rate parameter is interpreted as a data-simulation scale factor for the jet→tau fake rate.

3. The normalization of the $\mathrm{Z} \to \mu\mu$ background contributing to the $\tau_\mu\tau_{\text{1-prong}}$ selected sample is also allowed to float freely in the fit. This rate parameter accounts for the data-simulation scale factor associated with the efficiency of the extra-muon veto. This scale factor and its uncertainty are poorly known prior to the fit. The selection efficiency of loosely identified muons is 98-99% for muons stemming from Z decays. The uncertainty of O(1%) in the selection efficiency translates into an uncertainty of O(100%) in the muon veto efficiency. As a consequence, veto efficiency in data may largely differ from that in the simulated samples. Prefit distributions of the muon-track invariant mass for track $p_{\mathrm{T}} > 20$ GeV show a sizable excess in data compared to MC in the region under the Z peak, suggesting that the muon rejection efficiency in simulated events is a factor 2-3 higher than in data. Therefore, a freely floating rate parameter is introduced to control the contribution from processes with two genuine muons, where one muon is identified as a one-prong tau candidate.

4. One rate parameter is applied to Drell-Yan templates. Constraint on this parameter obtained in the $\mathrm{Z} \to \mu\mu$ control region is transferred to simulated Drell-Yan templates in the selected samples of $\tau_\mu\tau_{\text{1-prong}}$ candidates.

The following uncertainties are taken into account in the fit:

- Muon ID, isolation, and trigger efficiency (3%);
- Luminosity (1.5-2.0% depending on the data taking period);
- Uncertainties in the cross-sections of top pair and single-top production (7%), W+jets production (5%), and diboson processes (10%);
- Uncertainties in the normalization of the QCD multijet background: this uncertainty is estimated by varying the non-QCD background subtracted from data in the same-sign sideband by ±20%. This results in variations in the yield of QCD multijet events ranging from 5% to 40% depending on one-prong tau p_{T} bin and m_{T} region.

The uncertainties associated with the Drell-Yan cross-section are absorbed into the floating rate parameter, controlling the normalization of this process. Uncertainties in the jet energy scale and the unclustered energy scale have minimal impacts on the measurement and hence, are not considered.

The $m_{\mu,\mathrm{trk}}$ distributions in the signal region for the various p_{T} region, following the likelihood fit are shown in Fig. F.2, for the year 2018.

The scale factors extracted from the measurement are presented as a function of the track transverse momentum in Figs. F.3-F.5, for the years 2016–2018. As no dependence of the scale factors is observed with respect to the track p_{T}, the scale factors are fitted with a constant function. A flat scale factor is thus applied to all the events in the simulated samples used in the main analysis. To be conservative, p_{T} dependent uncertainties are modeled using a linear fit and incorporated to refine the flat scale factors applied to the simulated events. These uncertainties are represented by the blue bands in the plots and show the envelope of systematic variations of the slope of the linear function and the offset parameter.

F.4 Dependence of Scale Factors on $\Delta R_{\mathrm{iso,trk}}$

In the measurement discussed in the previous section, the isolation cone is centered on the track momentum. In the main analysis, however, the isolation cone is centered on the momentum of the muon originating from another tau decay in the same $\mathrm{a}_1 \to \tau\tau$ decay leg. To verify the applicability of the derived scale factors in the main analysis, the dependence of the scale factor on the angular separation between the direction of the isolation cone and the momentum of the track, $\Delta R_{\mathrm{iso,trk}}$, is studied. Four scenarios are considered: $\Delta R_{\mathrm{iso,trk}} = 0$, 0.15, 0.3, 0.45. The selection process used is the same as the one described previously, with one exception: the cut on the track momentum is increased to 10 GeV. This adjustment helps suppress background processes where one-prong tau candidates are mimicked by hadronic jets. As a result, this enables a more reliable and accurate measurement.

In these studies, for each event, the direction of the cone is generated randomly at a fixed value of $\Delta R_{\mathrm{iso,trk}}$ as

$$\begin{aligned} \eta_{\mathrm{iso}} &= \eta_{\mathrm{trk}} + \Delta R_{\mathrm{iso,trk}} \cdot \cos\phi^*, \\ \phi_{\mathrm{iso}} &= \phi_{\mathrm{trk}} + \Delta R_{\mathrm{iso,trk}} \cdot \sin\phi^*, \end{aligned} \tag{F.1}$$

where,

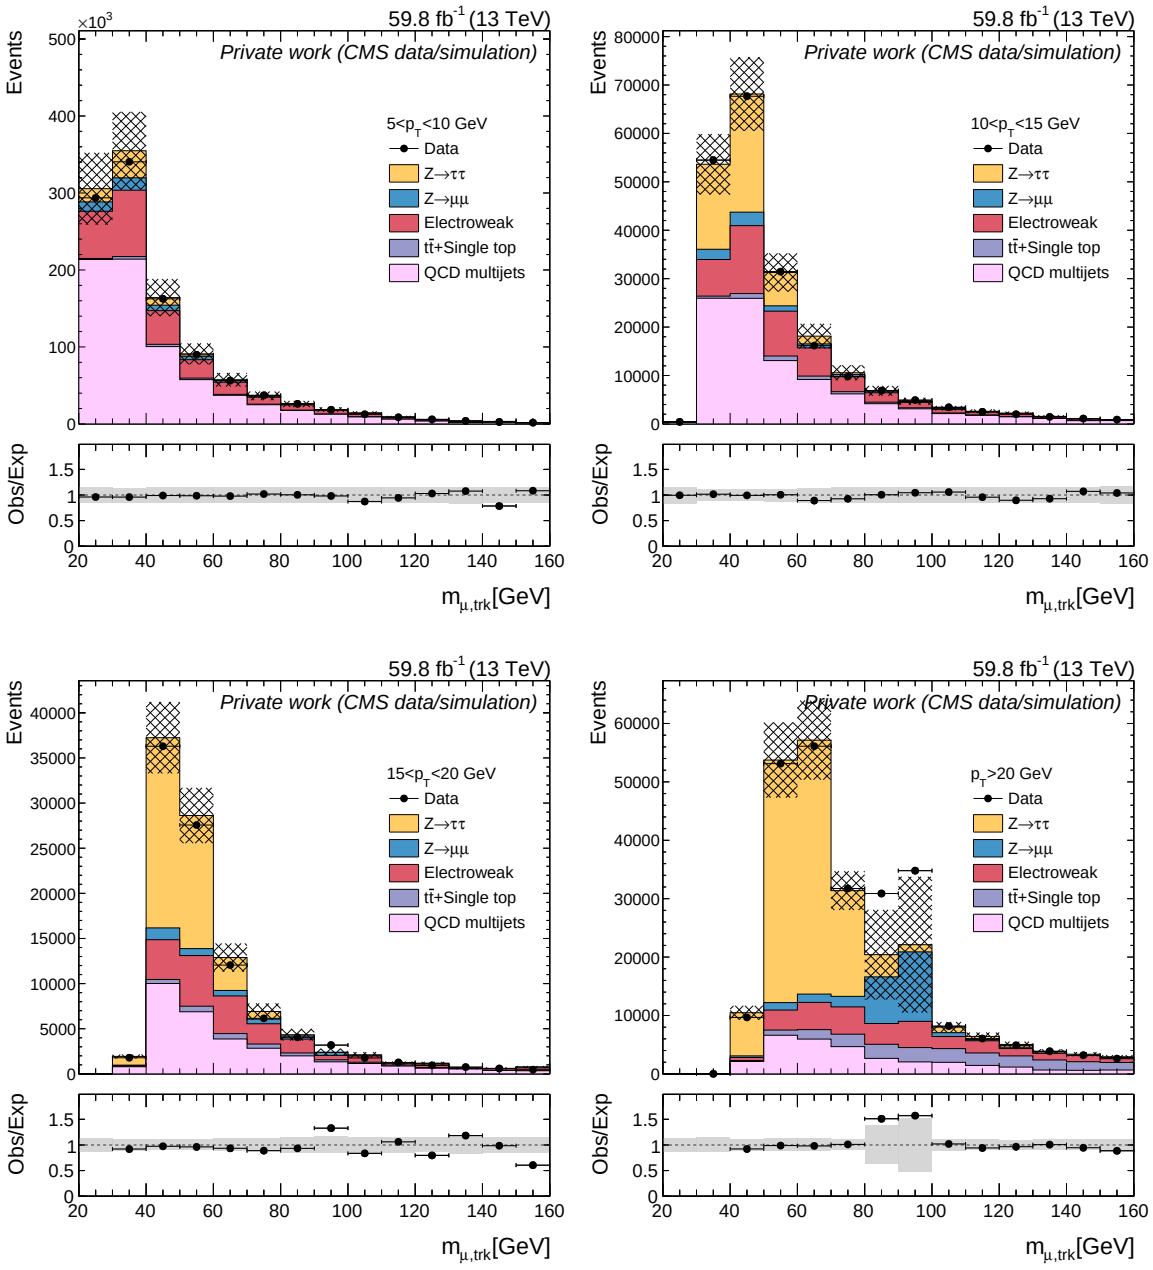

Figure F.1: Prefit distributions of $m_{\mu,\mathrm{trk}}$ in the selected sample of $\tau_\mu\tau_{\text{1-prong}}$ events in 2018 for the four different ranges of the track p_T: $5 < p_\mathrm{T} < 10\,\mathrm{GeV}$ (top left), $10 < p_\mathrm{T} < 15\,\mathrm{GeV}$ (top right), $15 < p_\mathrm{T} < 20\,\mathrm{GeV}$ (bottom left), and $p_\mathrm{T} > 20\,\mathrm{GeV}$ (bottom right).

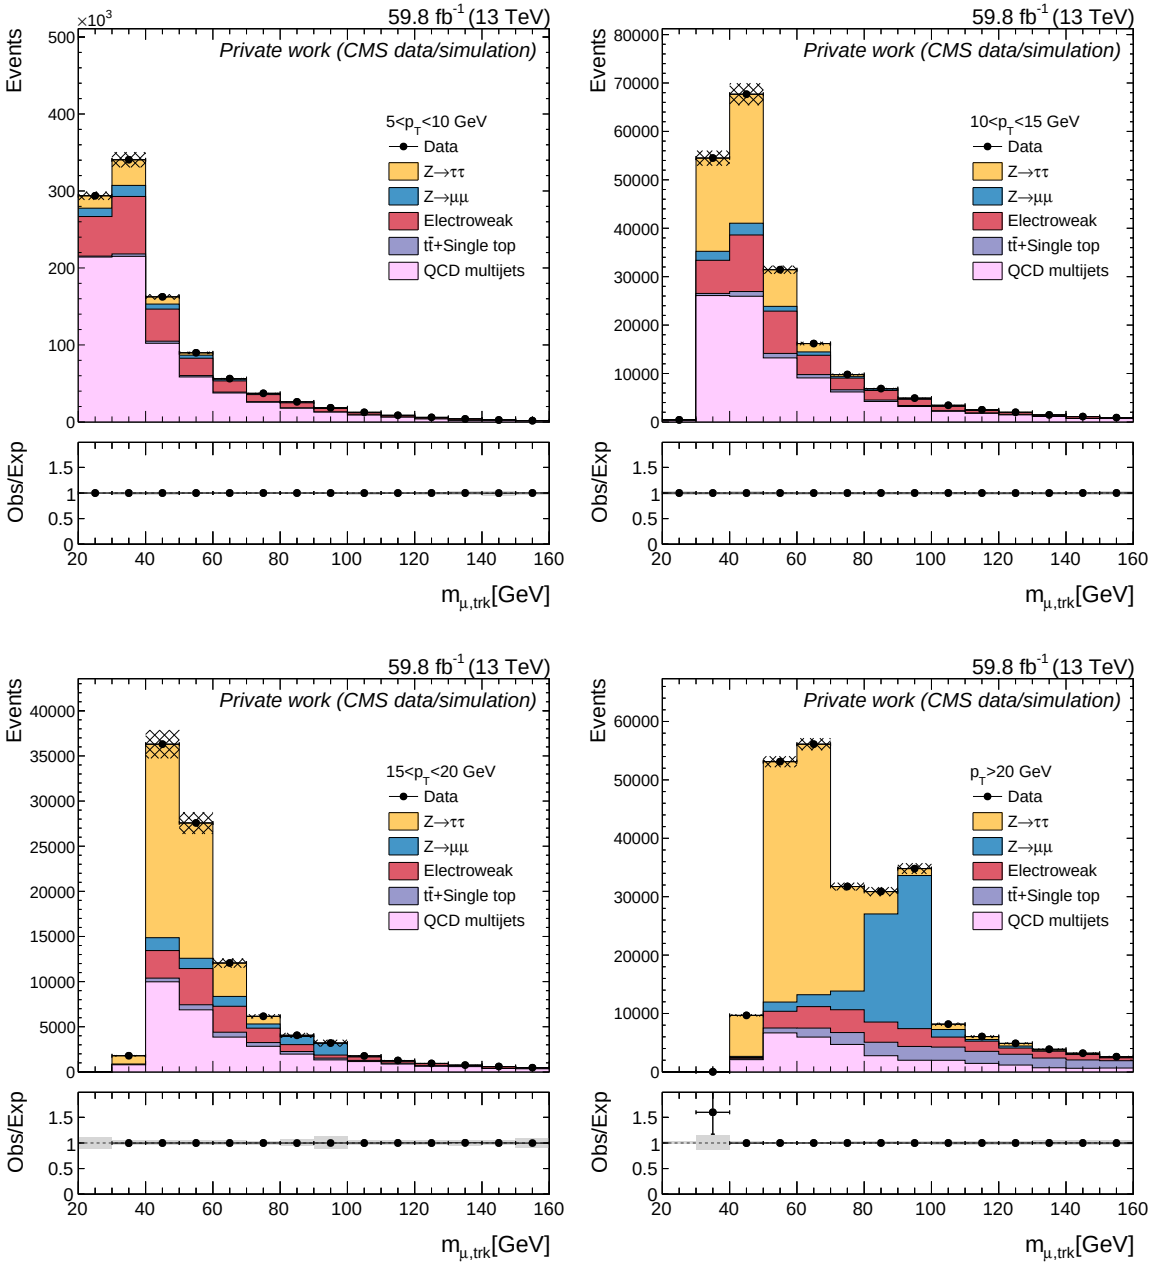

Figure F.2: Postfit distributions of $m_{\mu,\mathrm{trk}}$ in the selected sample of $\tau_\mu\tau_{\text{1-prong}}$ in 2018 for the four different ranges of the track p_T: $5 < p_\mathrm{T} < 10\,\mathrm{GeV}$ (top left), $10 < p_\mathrm{T} < 15\,\mathrm{GeV}$ (top right), $15 < p_\mathrm{T} < 20\,\mathrm{GeV}$ (bottom left), and $p_\mathrm{T} > 20\,\mathrm{GeV}$ (bottom right).

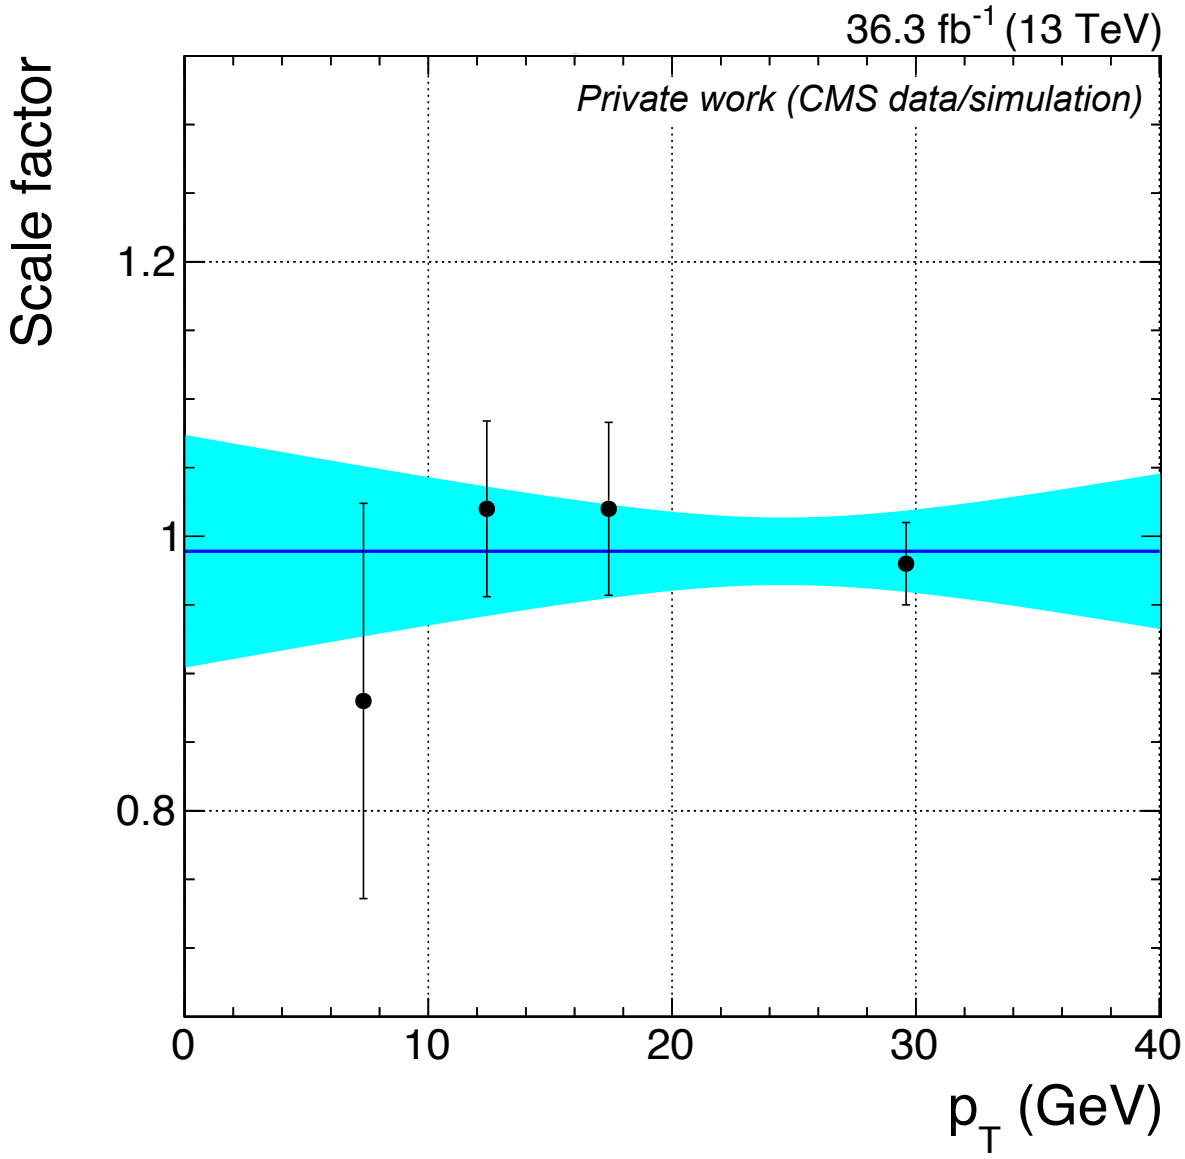

Figure F.3: The scale factor for the track isolation and identification efficiency as a function of the track p_{T} for 2016. As no dependence of the scale factors is observed with respect to the track p_{T}, the scale factors are fitted with a constant function (dark blue line). The blue band represents the uncertainty associated with the scale factor at each p_T value, modeled using a linear fit.

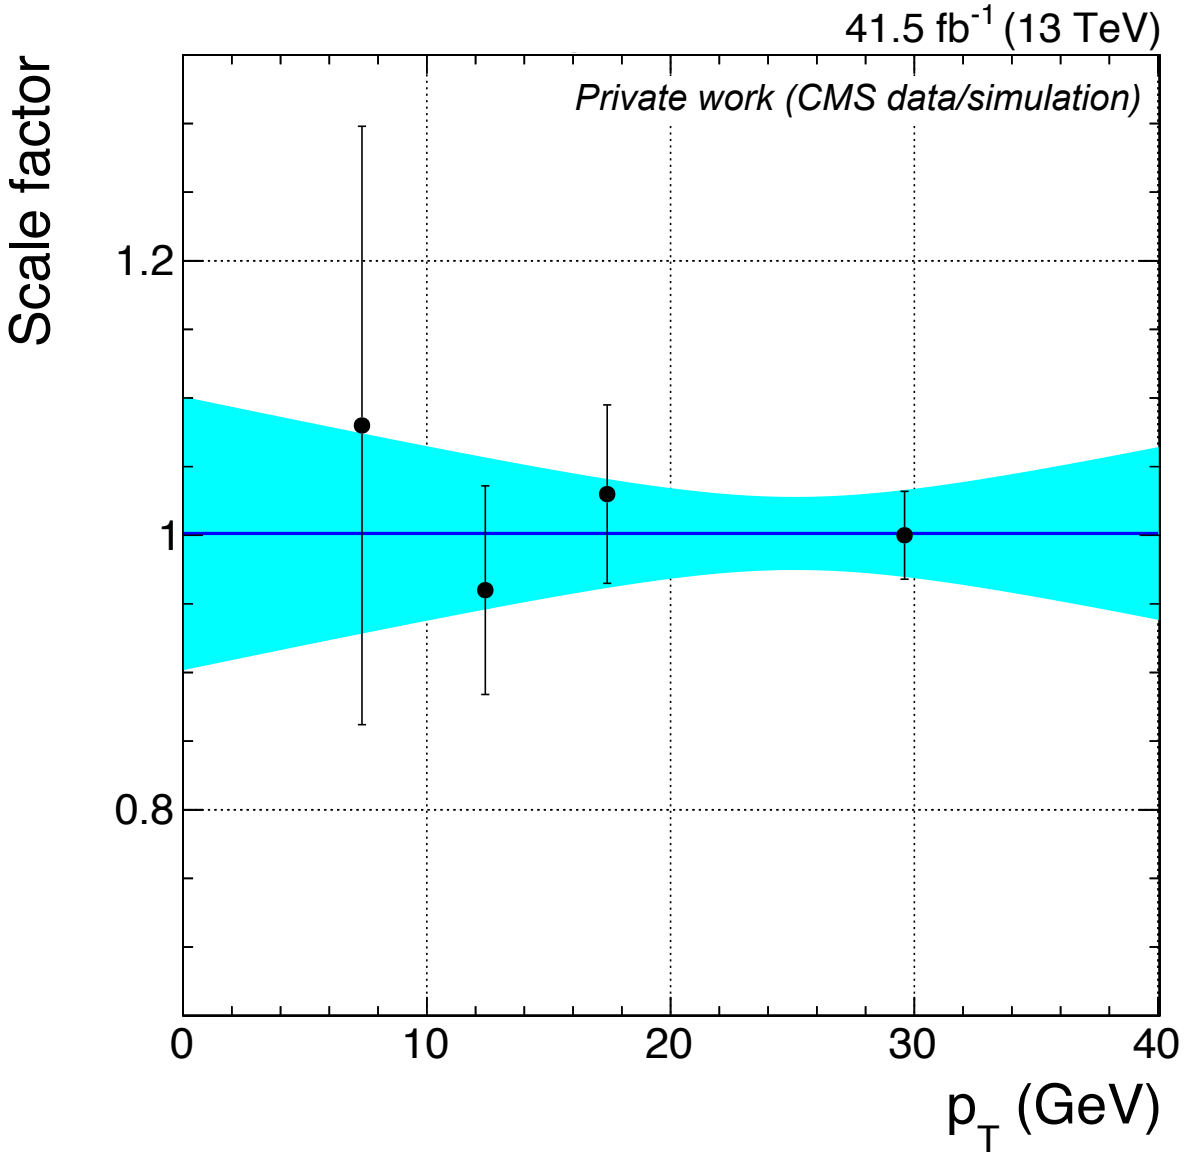

Figure F.4: The scale factor for the track isolation and identification efficiency as a function of the track p_{T} for 2017. As no dependence of the scale factors is observed with respect to the track p_{T}, the scale factors are fitted with a constant function (dark blue line). The blue band represents the uncertainty associated with the scale factor at each p_T value, modeled using a linear fit.

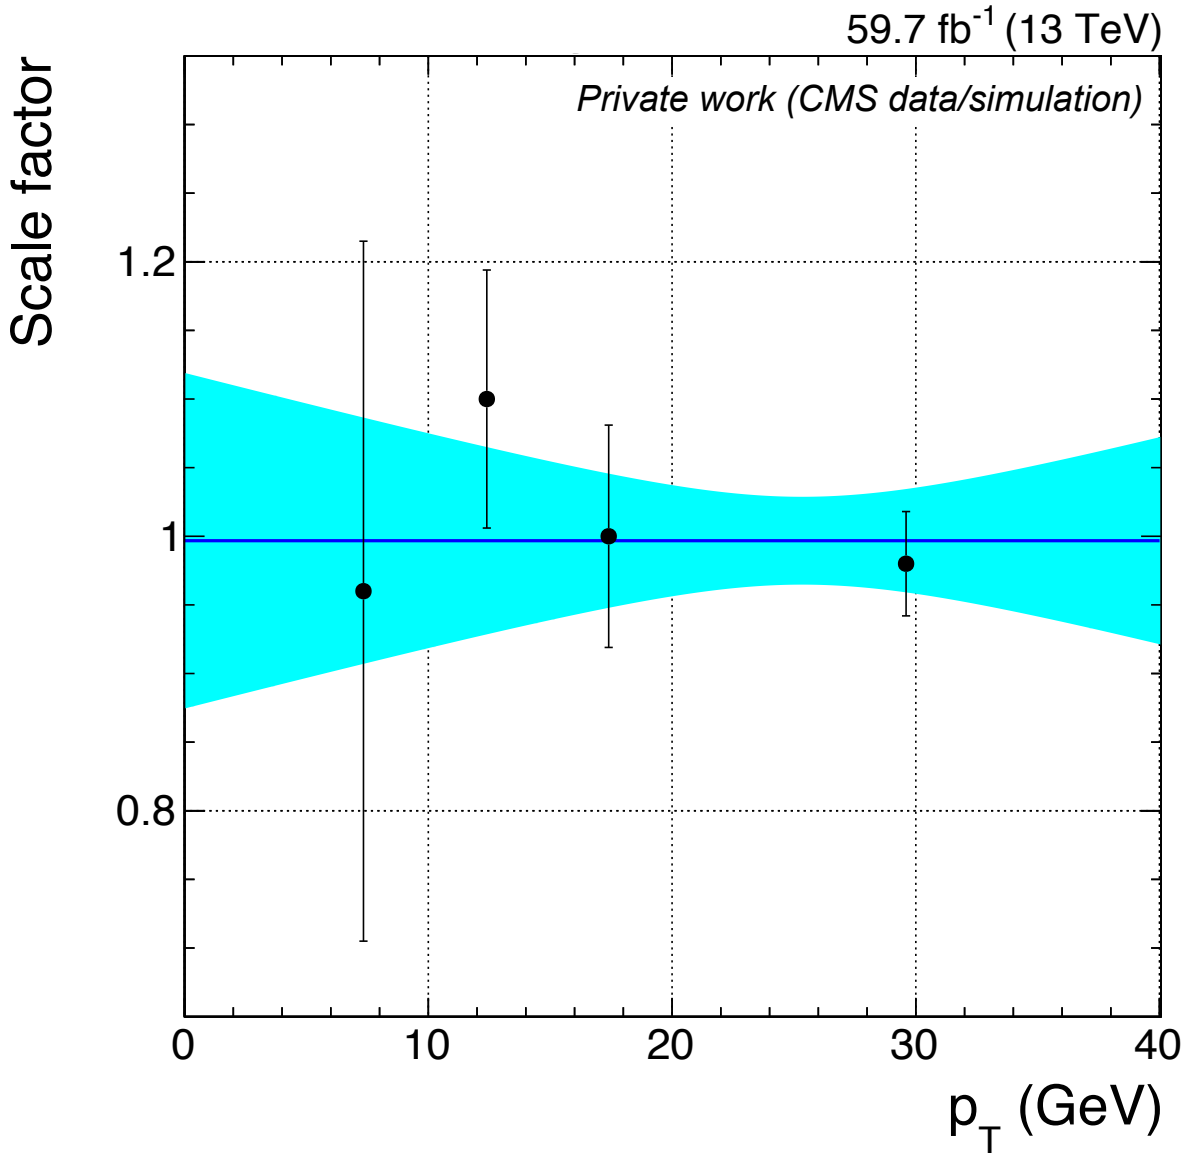

Figure F.5: The scale factor for the track isolation and identification efficiency as a function of the track p_{T} for 2018. As no dependence of the scale factors is observed with respect to the track p_{T}, the scale factors are fitted with a constant function (dark blue line). The blue band represents the uncertainty associated with the scale factor at each p_T value, modeled using a linear fit.

- η_{iso} is the pseudorapidity of the isolation cone;
- ϕ_{iso} is the azimuthal angle of the isolation cone;
- η_{trk} is pseudorapidity of the track;
- ϕ_{trk} is the azimuthal angle of the track;
- ϕ^* is randomly generated and uniformly distributed between 0 and 2π angle in the (η, ϕ) plane.

Figure F.6 shows the prefit $m_{\mu,\text{trk}}$ distributions for different values of $\Delta R_{\text{iso,trk}}$ for the year 2018, as an example. The corresponding postfit distributions are shown in Fig. F.7.

The scale factors extracted from the likelihood fit in these $\Delta R_{\text{iso,trk}}$ regions for the years 2016–2018 are shown in Figs. F.8–F.10. No apparent dependence of the scale factor on $\Delta R_{\text{iso,trk}}$ is observed, justifying the application of the scale factors determined at $\Delta R_{\text{iso,trk}} = 0$ in the main analysis.

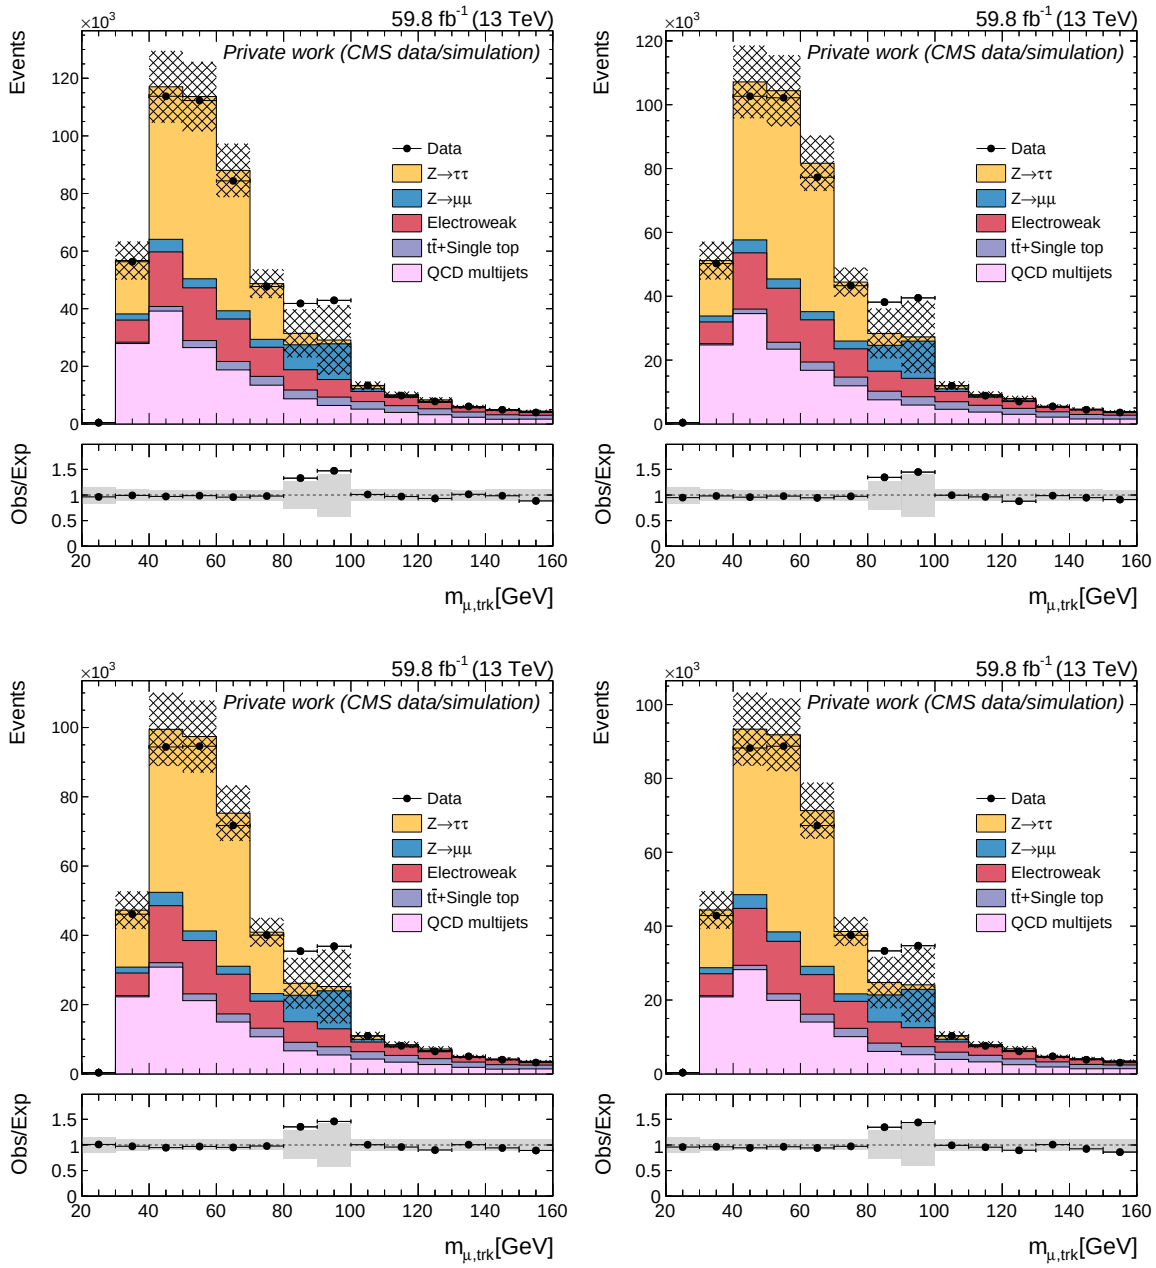

Figure F.6: Prefit distributions of $m_{\mu,\text{trk}}$ in the selected sample of $\tau_\mu\tau_{\text{1-prong}}$ in 2018 for the four different values of $\Delta R_{\text{iso,trk}} =:$ 0 (top left), 0.15 (top right), 0.30 (bottom left), and 0.45 (bottom right).

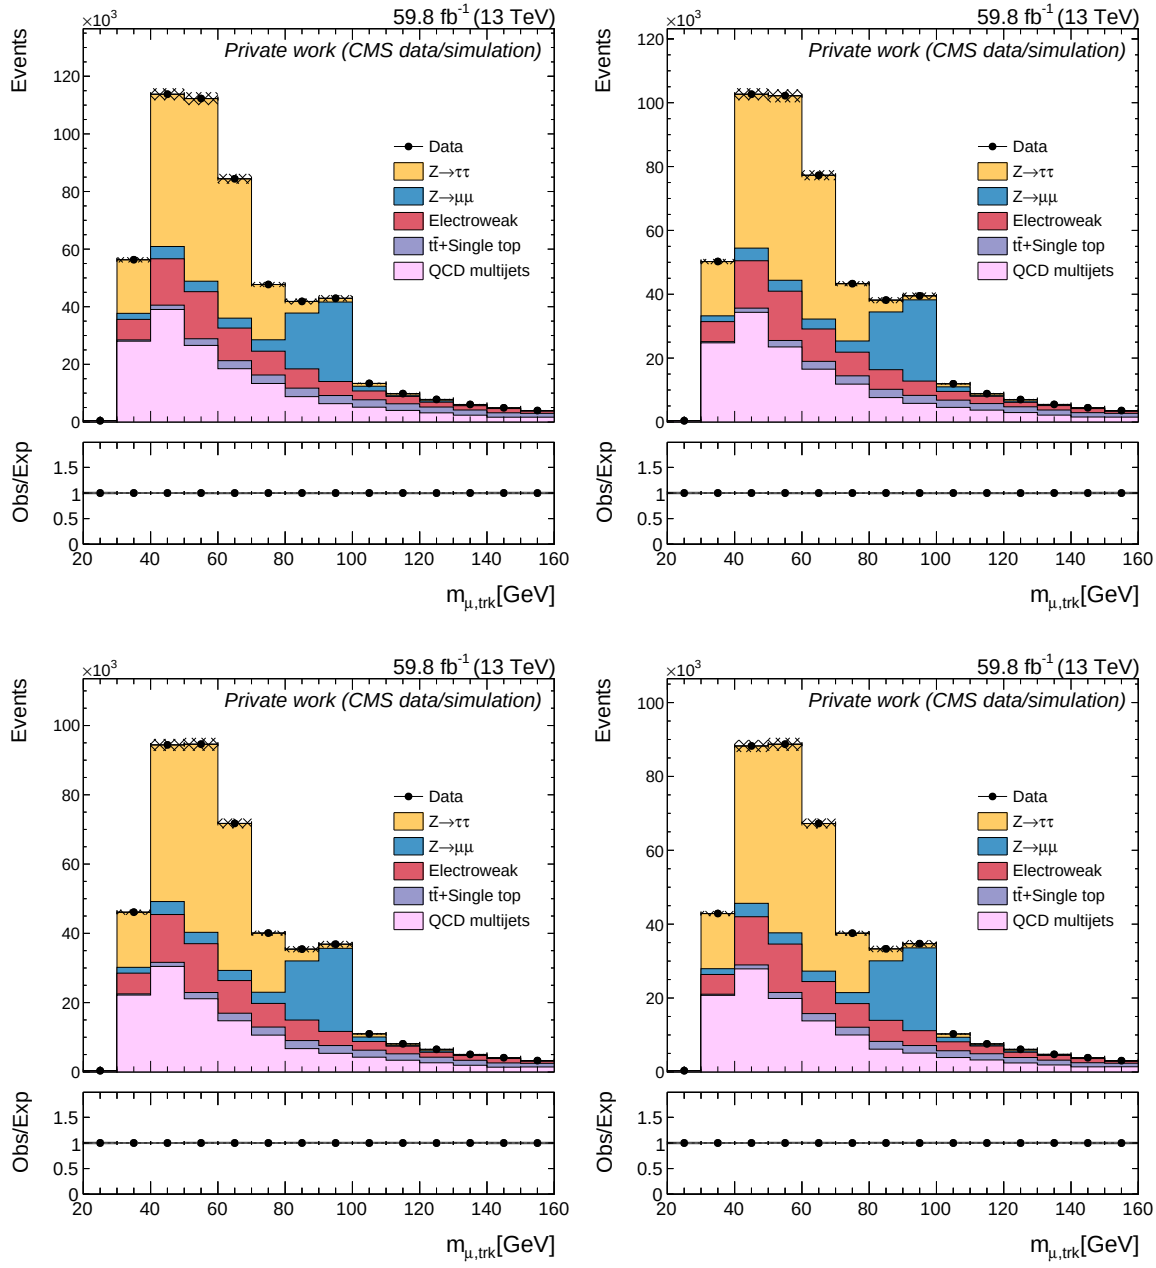

Figure F.7: Postfit distributions of $m_{\mu,\mathrm{trk}}$ in the selected sample of $\tau_\mu\tau_{\text{1-prong}}$ in 2018 for the four different values of $\Delta R_{\mathrm{iso,trk}} =:$ 0 (top left), 0.15 (top right), 0.30 (bottom left), and 0.45 (bottom right).

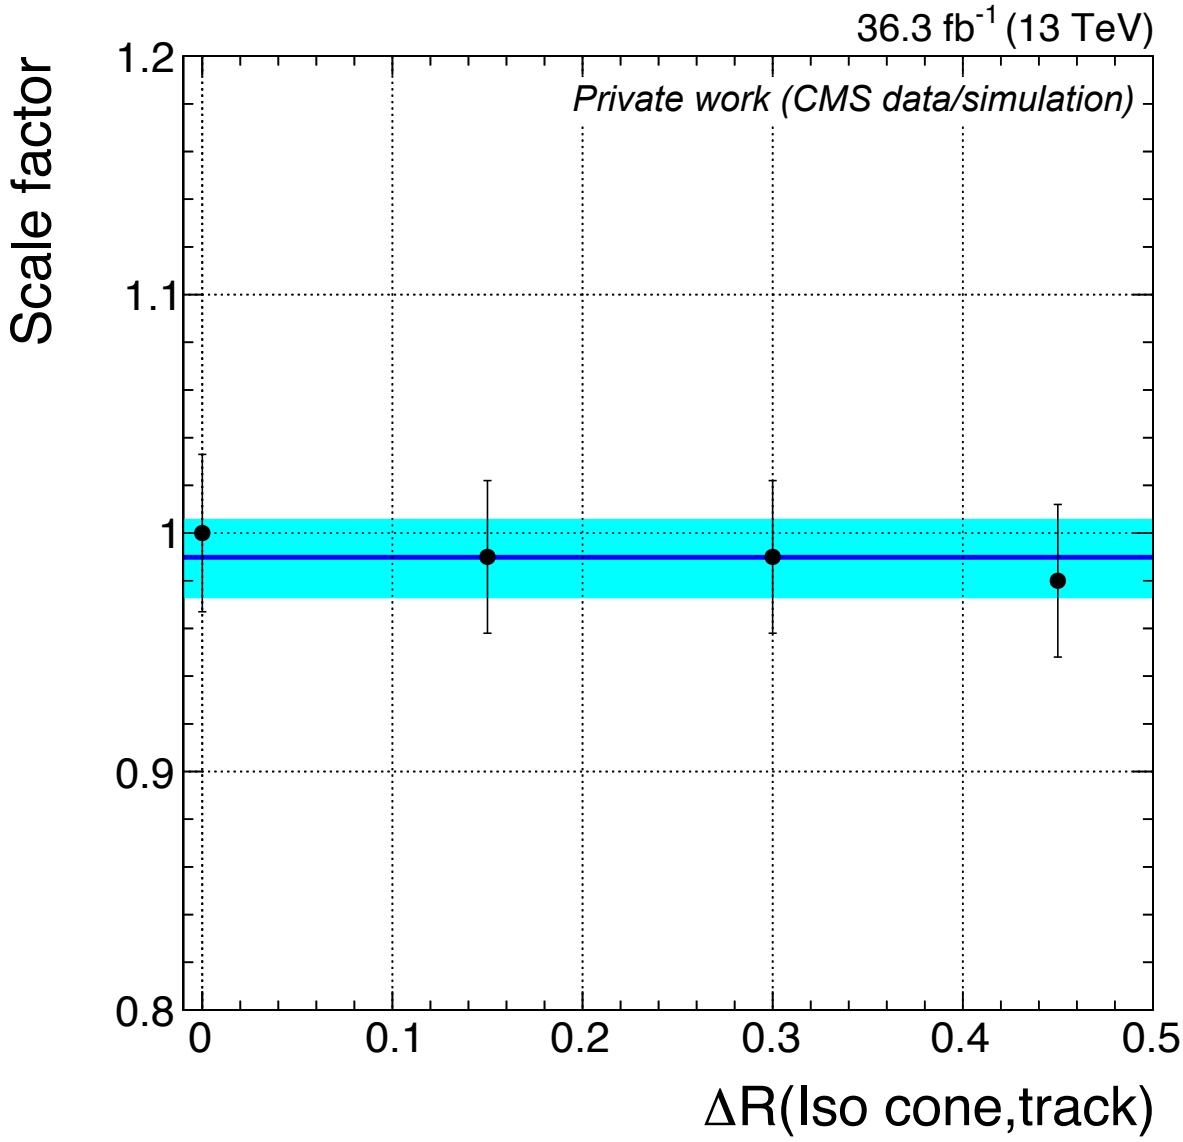

Figure F.8: Dependence of the scale factor for the track isolation and identification criteria on the angular separation between the direction of the isolation cone and track momentum in 2016.

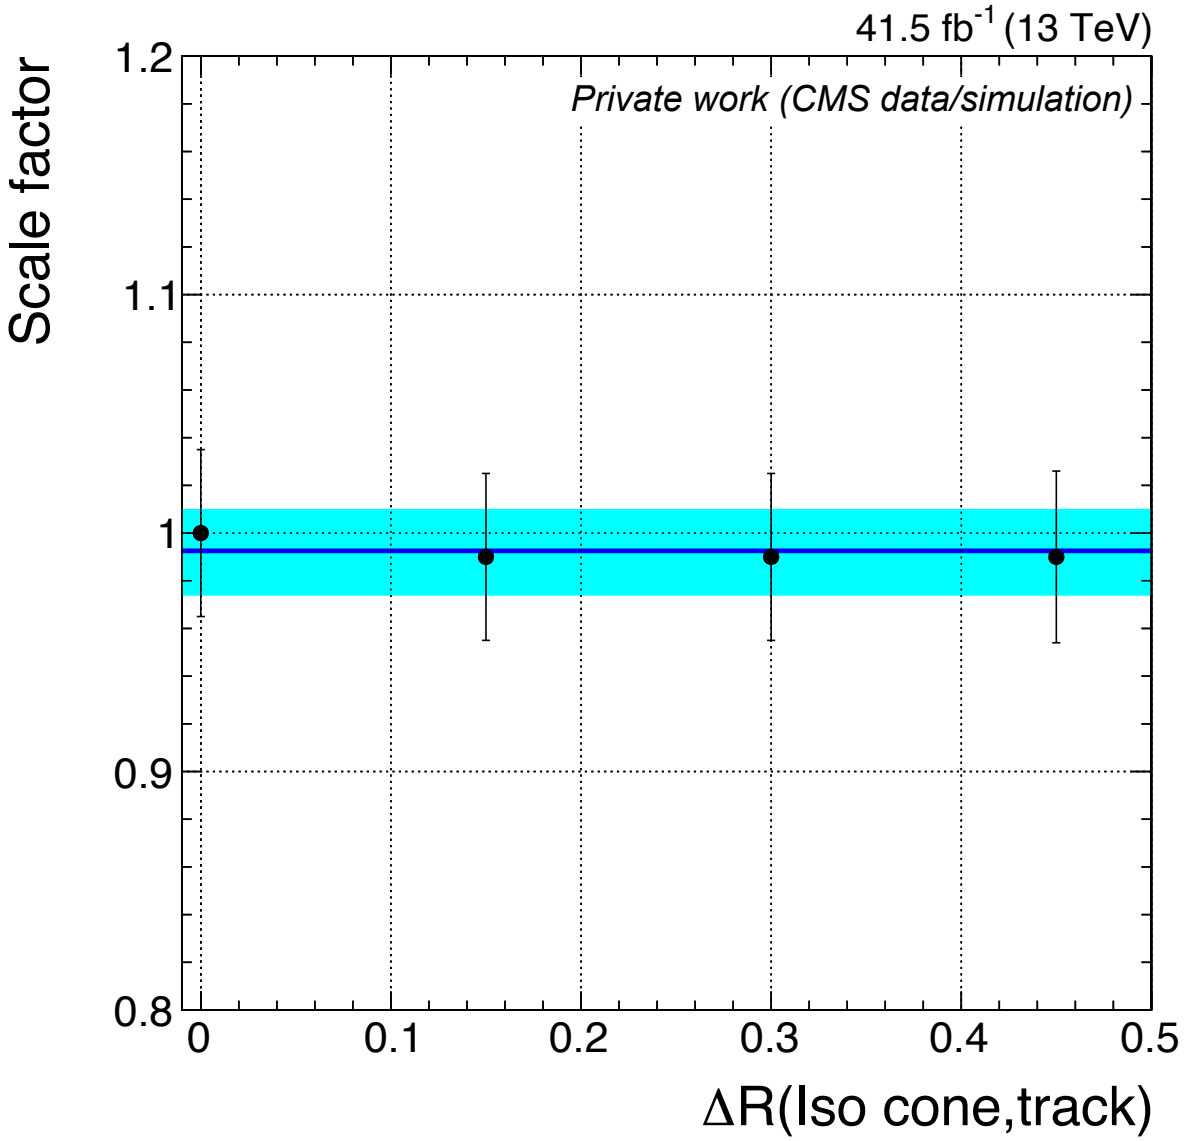

Figure F.9: Dependence of the scale factor for the track isolation and identification criteria on the angular separation between the direction of the isolation cone and track momentum in 2017.

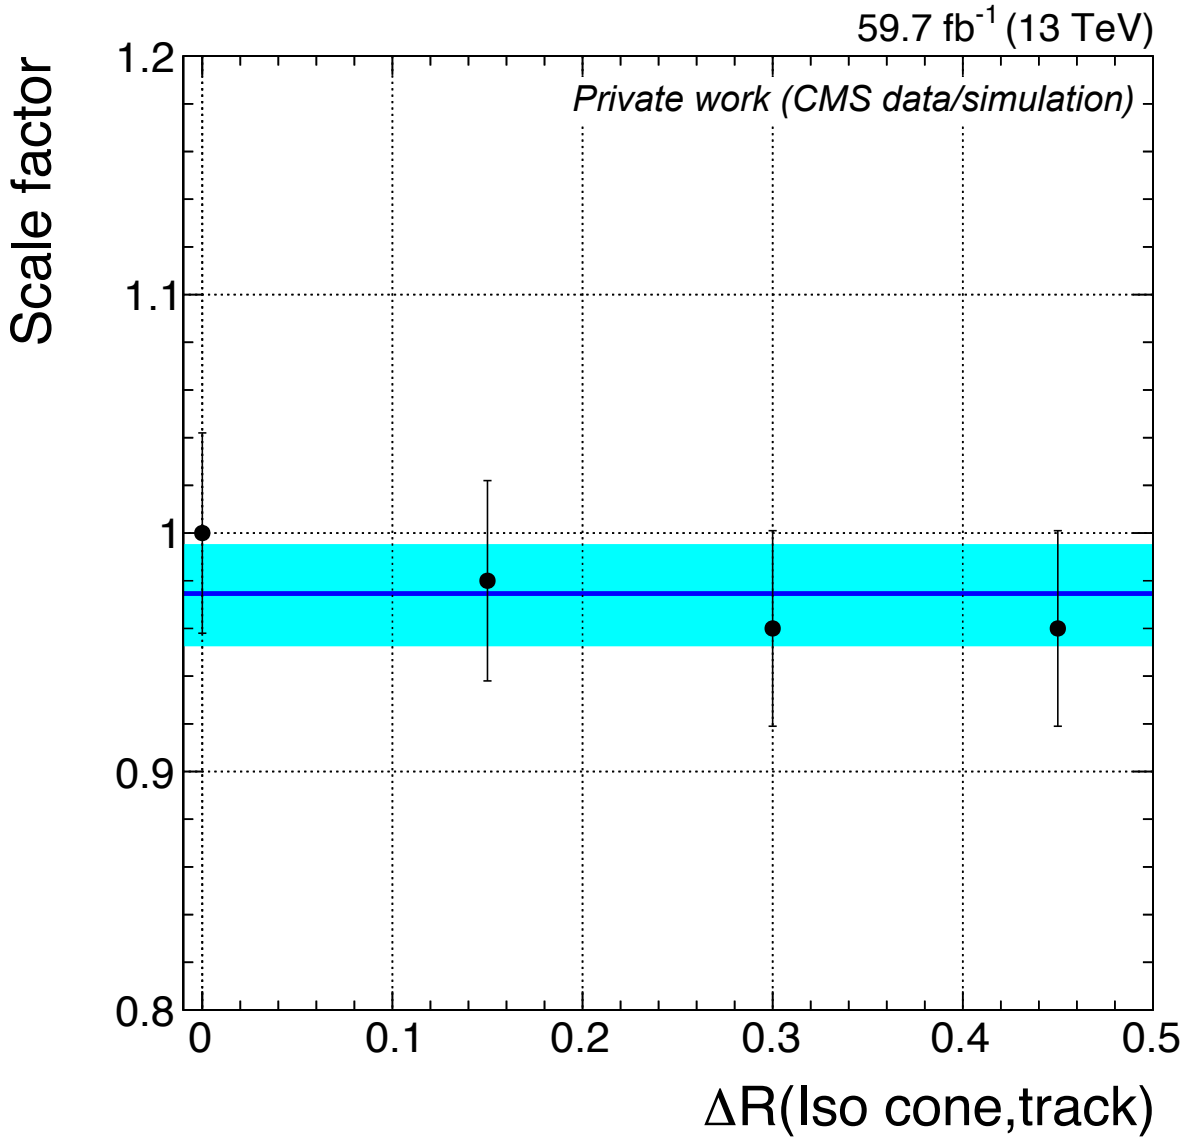

Figure F.10: Dependence of the scale factor for the track isolation and identification criteria on the angular separation between the direction of the isolation cone and track momentum in 2018.

Appendix G

Model-independent Upper Limits

Mass [GeV]	-2σ	-1σ	Expected	$+1\sigma$	$+2\sigma$	Observed
4	0.071	0.095	0.135	0.195	0.275	0.109
5	0.045	0.060	0.086	0.126	0.180	0.059
6	0.035	0.048	0.068	0.100	0.142	0.113
7	0.026	0.035	0.051	0.076	0.111	0.052
8	0.017	0.023	0.034	0.052	0.076	0.029
9	0.013	0.019	0.027	0.041	0.061	0.025
10	0.011	0.015	0.023	0.036	0.054	0.021
11	0.010	0.014	0.021	0.033	0.050	0.020
12	0.011	0.016	0.024	0.038	0.059	0.023
13	0.014	0.020	0.031	0.049	0.076	0.034
14	0.024	0.034	0.051	0.083	0.131	0.050
15	0.037	0.053	0.082	0.133	0.215	0.089

Table G.1: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM, for the 2016 analysis.

Mass [GeV]	-2σ	-1σ	Expected	$+1\sigma$	$+2\sigma$	Observed
4	0.062	0.084	0.118	0.172	0.243	0.152
5	0.046	0.062	0.088	0.129	0.183	0.067
6	0.040	0.055	0.078	0.114	0.163	0.037
7	0.028	0.038	0.054	0.081	0.117	0.025
8	0.021	0.029	0.042	0.062	0.091	0.021
9	0.016	0.021	0.031	0.047	0.069	0.019
10	0.013	0.019	0.027	0.041	0.061	0.019
11	0.013	0.019	0.028	0.042	0.063	0.021
12	0.015	0.022	0.032	0.049	0.073	0.027
13	0.020	0.027	0.040	0.062	0.093	0.037
14	0.031	0.043	0.064	0.099	0.150	0.068
15	0.063	0.087	0.129	0.200	0.307	0.120

Table G.2: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM, for the 2017 analysis.

Mass [GeV]	-2σ	-1σ	Expected	$+1\sigma$	$+2\sigma$	Observed
4	0.055	0.074	0.106	0.155	0.222	0.127
5	0.037	0.050	0.073	0.108	0.156	0.057
6	0.028	0.038	0.056	0.083	0.120	0.052
7	0.019	0.026	0.038	0.056	0.083	0.034
8	0.014	0.020	0.029	0.044	0.065	0.027
9	0.011	0.015	0.022	0.034	0.050	0.018
10	0.009	0.013	0.018	0.029	0.043	0.014
11	0.009	0.012	0.018	0.028	0.043	0.013
12	0.011	0.015	0.023	0.037	0.056	0.016
13	0.010	0.014	0.022	0.035	0.055	0.015
14	0.019	0.027	0.041	0.066	0.105	0.028
15	0.038	0.054	0.085	0.140	0.227	0.054

Table G.3: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM, for the 2018 analysis.

Mass [GeV]	-2σ	-1σ	Expected	$+1\sigma$	$+2\sigma$	Observed
4	0.035	0.048	0.067	0.095	0.131	0.077
5	0.024	0.032	0.045	0.065	0.090	0.028
6	0.019	0.026	0.036	0.052	0.073	0.031
7	0.013	0.018	0.025	0.036	0.051	0.016
8	0.009	0.013	0.018	0.027	0.037	0.012
9	0.007	0.010	0.014	0.020	0.029	0.009
10	0.006	0.008	0.012	0.017	0.025	0.008
11	0.005	0.008	0.011	0.017	0.024	0.008
12	0.007	0.009	0.013	0.020	0.029	0.010
13	0.008	0.010	0.014	0.022	0.032	0.012
14	0.012	0.017	0.025	0.038	0.056	0.020
15	0.023	0.032	0.046	0.071	0.106	0.037

Table G.4: The observed and expected upper limits at 95% confidence level on the product of the signal cross section and the branching fraction $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{a_1 a_1})\mathcal{B}^2(\mathrm{a_1} \to \tau\tau)$, relative to the inclusive Higgs boson production cross section σ_{SM} predicted in the SM, for the Run 2 combination.

Bibliography

[1] UA1 Collaboration, "Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV", *Phys. Lett. B* **122** (1983) 103–116, doi:10.1016/0370-2693(83)91177-2. (Cited on page 1)

[2] UA2 Collaboration, "Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider", *Phys. Lett. B* **122** (1983) 476–485, doi:10.1016/0370-2693(83)91605-2. (Cited on page 1)

[3] E598 Collaboration, "Experimental Observation of a Heavy Particle J", *Phys. Rev. Lett.* **33** (1974) 1404–1406, doi:10.1103/PhysRevLett.33.1404. (Cited on page 1)

[4] SLAC-SP-017 Collaboration, "Discovery of a Narrow Resonance in e^+e^- Annihilation", *Phys. Rev. Lett.* **33** (1974) 1406–1408, doi:10.1103/PhysRevLett.33.1406. (Cited on page 1)

[5] CDF Collaboration, "Observation of top quark production in $\bar{p}p$ collisions", *Phys. Rev. Lett.* **74** (1995) 2626–2631, doi:10.1103/PhysRevLett.74.2626, arXiv:hep-ex/9503002. (Cited on page 1)

[6] ATLAS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC", *Phys. Lett. B* **716** (2012) 1–29, doi:10.1016/j.physletb.2012.08.020, arXiv:1207.7214. (Cited on pages 1, 18, and 38)

[7] CMS Collaboration, "Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC", *Phys. Lett. B* **716** (2012) 30–61, doi:10.1016/j.physletb.2012.08.021, arXiv:1207.7235. (Cited on pages 1, 18, and 38)

[8] CMS Collaboration, "Observation of a New Boson with Mass Near 125 GeV in pp Collisions at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **06** (2013) 081, doi:10.1007/JHEP06(2013)081, arXiv:1303.4571. (Cited on pages 1, 18, and 38)

[9] CMS Collaboration, "A portrait of the Higgs boson by the CMS experiment ten years after the discovery.", *Nature* **607** (2022) 60, doi:10.1038/s41586-022-04892-x, arXiv:2207.00043. (Cited on pages 1, 20, 21, 22, 82, 94, 150, 151, and 152)

[10] ATLAS Collaboration, "A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery", *Nature* **607** (2022) 52,

doi:10.1038/s41586-022-04893-w, arXiv:2207.00092. [Erratum: Nature 612, E24 (2022)]. (Cited on pages 1, 20, 22, and 82)

[11] D. Curtin et al., "Exotic decays of the 125 GeV Higgs boson", *Phys. Rev. D* **90** (2014), no. 7, 075004, doi:10.1103/PhysRevD.90.075004, arXiv:1312.4992. (Cited on pages 2, 23, and 28)

[12] CMS Collaboration, "CMS Tracker Alignment Preliminary Results for 2021 startup with 900 GeV collision data", https://twiki.cern.ch/twiki/bin/view/CMSPublic/TkAlPreliminaryPilotTestBeamResultsForLHCC. Accessed: Jan. 28, 2025. (Cited on pages 3, 48, and 167)

[13] CMS Collaboration, "Tracker alignment performance in 2021", CMS Detector Performance Summary, CMS-DP-2022-017, 2022. (Cited on pages 3, 48, 167, and 169)

[14] CMS Collaboration, "Tracker Alignment Performance in 2022", CMS Detector Performance Summary, CMS-DP-2022-044, 2022. (Cited on pages 3, 48, and 167)

[15] CMS Collaboration, "Tracker alignment performance in 2022 (addendum)", CMS Detector Performance Summary, CMS-DP-2024-078, 2022. (Cited on pages 3, 48, and 167)

[16] CMS Collaboration, "Tracker alignment performance in early 2023", CMS Detector Performance Summary, CMS-DP-2023-039, 2023. (Cited on pages 3, 48, 167, and 170)

[17] CMS Collaboration, "CMS tracker alignment performance results for Run 3 reprocessing", CMS Detector Performance Summary, CMS-DP-2024-071, 2024. (Cited on pages 3, 48, 167, 172, 173, and 174)

[18] CMS Collaboration, "Search for light pseudoscalar bosons produced in Higgs boson decays in the 4τ and $2\mu 2\tau$ final states in proton-proton collisions at $\sqrt{s} = 13$ TeV", Technical Report, CMS-PAS-SUS-24-002, CERN, Geneva, 2024. (Cited on pages 3, 82, and 156)

[19] D. Galbraith and C. Burgard, "Standard model of physics", https://texample.net/model-physics/. Accessed: Jan. 31, 2025. (Cited on page 7)

[20] Particle Data Group Collaboration, "Review of particle physics", *Phys. Rev. D* **110** (2024), no. 3, 030001, doi:10.1103/PhysRevD.110.030001. (Cited on pages 7, 8, 10, 13, 18, 19, 77, and 78)

[21] E. Noether, "Invariant Variation Problems", *Gott. Nachr.* **1918** (1918) 235–257, doi:10.1080/00411457108231446, arXiv:physics/0503066. (Cited on page 7)

[22] M. E. Peskin and D. V. Schroeder, "An Introduction to quantum field theory". Addison-Wesley, Reading, USA, 1995, doi:10.1201/9780429503559. ISBN: 978-0-201-50397-5, 978-0-429-50355-9, 978-0-429-49417-8. (Cited on page 8)

[23] D. Griffiths, "Introduction to elementary particles". Wiley-VCH, 2008. ISBN: 978-3-527-40601-2. (Cited on pages 8, 9, and 17)

[24] R. K. Ellis, W. J. Stirling, and B. R. Webber, "QCD and collider physics", volume 8. Cambridge University Press, 2, 2011, doi:10.1017/CBO9780511628788. ISBN: 978-0-511-82328-2, 978-0-521-54589-1. (Cited on pages 9 and 60)

[25] D. J. Gross and F. Wilczek, "Ultraviolet Behavior of Nonabelian Gauge Theories", *Phys. Rev. Lett.* **30** (1973) 1343–1346, doi:10.1103/PhysRevLett.30.1343. (Cited on page 10)

[26] H. D. Politzer, "Reliable Perturbative Results for Strong Interactions?", *Phys. Rev. Lett.* **30** (1973) 1346–1349, doi:10.1103/PhysRevLett.30.1346. (Cited on page 10)

[27] J. Greensite, "An introduction to the confinement problem", volume 821. Springer, 2011, doi:10.1007/978-3-642-14382-3. (Cited on page 11)

[28] S. L. Glashow, "Partial Symmetries of Weak Interactions", *Nucl. Phys.* **22** (1961) 579–588, doi:10.1016/0029-5582(61)90469-2. (Cited on page 11)

[29] S. Weinberg, "A Model of Leptons", *Phys. Rev. Lett.* **19** (1967) 1264–1266, doi:10.1103/PhysRevLett.19.1264. (Cited on pages 11 and 17)

[30] A. Salam and J. C. Ward, "Weak and electromagnetic interactions", *Nuovo Cim.* **11** (1959) 568–577, doi:10.1007/BF02726525. (Cited on page 11)

[31] F. Halzen and A. D. Martin, "Quarks and Leptones: An Introductory Course in Modern Particle Physics". Wiley & Sons Ltd, 1984. ISBN: 978-0-471-88741-6. (Cited on page 12)

[32] C. S. Wu et al., "Experimental Test of Parity Conservation in β Decay", *Phys. Rev.* **105** (1957) 1413–1414, doi:10.1103/PhysRev.105.1413. (Cited on page 12)

[33] F. Englert and R. Brout, "Broken Symmetry and the Mass of Gauge Vector Mesons", *Phys. Rev. Lett.* **13** (1964) 321–323, doi:10.1103/PhysRevLett.13.321. (Cited on page 14)

[34] P. W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons", *Phys. Rev. Lett.* **13** (1964) 508–509, doi:10.1103/PhysRevLett.13.508. (Cited on page 14)

[35] P. W. Higgs, "Broken symmetries, massless particles and gauge fields", *Phys. Lett.* **12** (1964) 132–133, doi:10.1016/0031-9163(64)91136-9. (Cited on page 14)

[36] P. W. Higgs, "Spontaneous Symmetry Breakdown without Massless Bosons", *Phys. Rev.* **145** (1966) 1156–1163, doi:10.1103/PhysRev.145.1156. (Cited on page 14)

[37] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, "Global Conservation Laws and Massless Particles", *Phys. Rev. Lett.* **13** (1964) 585–587, doi:10.1103/PhysRevLett.13.585. (Cited on page 14)

[38] J. Goldstone, "Field Theories with Superconductor Solutions", *Nuovo Cim.* **19** (1961) 154–164, doi:10.1007/BF02812722. (Cited on page 16)

[39] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, "Evidence for the 2π Decay of the K_2^0 Meson", *Phys. Rev. Lett.* **13** (1964) 138–140, doi:10.1103/PhysRevLett.13.138. (Cited on page 18)

[40] N. Cabibbo, "Unitary Symmetry and Leptonic Decays", *Phys. Rev. Lett.* **10** (1963) 531–533, doi:10.1103/PhysRevLett.10.531. (Cited on page 18)

[41] M. Kobayashi and T. Maskawa, "CP Violation in the Renormalizable Theory of Weak Interaction", *Prog. Theor. Phys.* **49** (1973) 652–657, doi:10.1143/PTP.49.652. (Cited on page 18)

[42] Planck Collaboration, "Planck 2018 results. VI. Cosmological parameters", *Astron. Astrophys.* **641** (2020) A6, doi:10.1051/0004-6361/201833910, arXiv:1807.06209. [Erratum: Astron.Astrophys. 652, C4 (2021)]. (Cited on page 18)

[43] KATRIN Collaboration, "Direct neutrino-mass measurement with sub-electronvolt sensitivity", *Nature Phys.* **18** (2022), no. 2, 160–166, doi:10.1038/s41567-021-01463-1, arXiv:2105.08533. (Cited on page 19)

[44] R. D. Peccei, "The Strong CP problem and axions", *Lect. Notes Phys.* **741** (2008) 3–17, doi:10.1007/978-3-540-73518-2_1, arXiv:hep-ph/0607268. (Cited on page 19)

[45] S. Weinberg, "Implications of Dynamical Symmetry Breaking", *Phys. Rev. D* **13** (1976) 974–996, doi:10.1103/PhysRevD.19.1277. [Addendum: Phys.Rev.D 19, 1277–1280 (1979)]. (Cited on page 19)

[46] LHC Higgs Cross Section Working Group Collaboration, D. de Florian et al., "Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector",. 10, 2016. doi:10.23731/CYRM-2017-002. (Cited on pages 21 and 99)

[47] CMS Collaboration, "Measurement of the Higgs boson production rate in association with top quarks in final states with electrons, muons, and hadronically decaying tau leptons at $\sqrt{s} = 13$ TeV", *Eur. Phys. J. C* **81** (2021), no. 4, 378, doi:10.1140/epjc/s10052-021-09014-x, arXiv:2011.03652. (Cited on page 22)

[48] CMS Collaboration, "Evidence for Higgs boson decay to a pair of muons", *JHEP* **01** (2021) 148, doi:10.1007/JHEP01(2021)148, arXiv:2009.04363. (Cited on pages 22 and 39)

[49] P. Fayet, "Supergauge Invariant Extension of the Higgs Mechanism and a Model for the electron and Its Neutrino", *Nucl. Phys. B* **90** (1975) 104–124, doi:10.1016/0550-3213(75)90636-7. (Cited on page 22)

[50] P. Fayet, "Spontaneously Broken Supersymmetric Theories of Weak, Electromagnetic and Strong Interactions", *Phys. Lett. B* **69** (1977) 489, doi:10.1016/0370-2693(77)90852-8. (Cited on page 22)

[51] U. Ellwanger, C. Hugonie, and A. M. Teixeira, "The next-to-minimal supersymmetric standard model", *Phys. Rept.* **496** (2010) 1, doi:10.1016/j.physrep.2010.07.001, arXiv:0910.1785. (Cited on page 23)

[52] M. Maniatis, "The next-to-minimal supersymmetric extension of the standard model reviewed", *Int. J. Mod. Phys. A* **25** (2010) 3505, doi:10.1142/S0217751X10049827, arXiv:0906.0777. (Cited on page 23)

[53] T. Robens and T. Stefaniak, "Status of the Higgs Singlet Extension of the Standard Model after LHC Run 1", *Eur. Phys. J. C* **75** (2015) 104, doi:10.1140/epjc/s10052-015-3323-y, arXiv:1501.02234. (Cited on page 23)

[54] M. Cepeda, S. Gori, V. M. Outschoorn, and J. Shelton, "Exotic Higgs Decays", doi:10.1146/annurev-nucl-102319-024147, arXiv:2111.12751. (Cited on page 23)

[55] S. Baum and N. R. Shah, "Two Higgs Doublets and a Complex Singlet: Disentangling the Decay Topologies and Associated Phenomenology", *JHEP* **12** (2018) 044, doi:10.1007/JHEP12(2018)044, arXiv:1808.02667. (Cited on pages 23 and 26)

[56] T. D. Lee, "A theory of spontaneous T violation", *Phys. Rev. D* **8** (1973) 1226, doi:10.1103/PhysRevD.8.1226. (Cited on page 23)

[57] G. C. Branco et al., "Theory and phenomenology of two-Higgs-doublet models", *Phys. Rept.* **516** (2012) 1, doi:10.1016/j.physrep.2012.02.002, arXiv:1106.0034. (Cited on page 23)

[58] T. Alanne, K. Kainulainen, K. Tuominen, and V. Vaskonen, "Baryogenesis in the two doublet and inert singlet extension of the standard model", *JCAP* **08** (2016) 057, doi:10.1088/1475-7516/2016/08/057, arXiv:1607.03303. (Cited on page 23)

[59] W. Su, A. G. Williams, and M. Zhang, "Strong first order electroweak phase transition in 2HDM confronting future Z & Higgs factories", *JHEP* **04** (2021) 219, doi:10.1007/JHEP04(2021)219, arXiv:2011.04540. (Cited on page 23)

[60] J. Dutta, G. Moortgat-Pick, and M. Schreiber, "Phenomenology of the dark matter sector in the 2HDM extended with complex scalar singlet", *Eur. Phys. J. Plus* **140** (2025), no. 1, 87, doi:10.1140/epjp/s13360-025-06015-7. (Cited on page 23)

[61] ATLAS, "Summary of Exotic Higgs Boson Decays from the ATLAS Experiment", https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-008/. Accessed: Jan. 31, 2025. (Cited on page 29)

[62] CMS, "Summary of 2HDM+S searches at 13 TeV from CMS", https://twiki.cern.ch/twiki/bin/view/CMSPublic/Summary2HDMSRun2. Accessed: Jan. 31, 2025. (Cited on pages 29, 30, and 31)

[63] L. Evans and P. Bryant, "LHC Machine", *JINST* **3** (2008) S08001, doi:10.1088/1748-0221/3/08/S08001. (Cited on pages 33 and 37)

[64] S. Myers and E. Picasso, "The Design, construction and commissioning of the CERN Large Electron Positron collider", *Contemp. Phys.* **31** (1990) 387–403, doi:10.1080/00107519008213789. (Cited on page 33)

[65] E. Boltezar et al., "Performance of the new CERN 50-MeV LINAC", *IEEE Trans. Nucl. Sci.* **26** (1979) 3674–3676, doi:10.1109/TNS.1979.4330576. (Cited on page 34)

[66] M. Vretenar et al., "Linac4 design report", volume 6/2020 of *CERN Yellow Reports: Monographs*. CERN, Geneva, 9, 2020, doi:10.23731/CYRM-2020-006. ISBN: 978-92-9083-579-0, 978-92-9083-580-6. (Cited on page 34)

[67] K. H. Reich, "The CERN Proton Synchrotron Booster", *IEEE Trans. Nucl. Sci.* **16** (1969) 959–961, doi:10.1109/TNS.1969.4325414. (Cited on page 34)

[68] D. Cundy and S. Gilardoni, "The Proton Synchrotron (PS): At the Core of the CERN Accelerators", *Adv. Ser. Direct. High Energy Phys.* **27** (2017) 39–85, doi:10.1142/9789814749145_0003. (Cited on page 34)

[69] J. B. Adams, "The CERN 400 GeV Proton Synchrotron (CERN SPS)", in *10th International Conference on High-Energy Accelerators*, volume 1. 1977, https://cds.cern.ch/record/2048898. (Cited on page 34)

[70] E. Mobs, "The CERN accelerator complex - August 2018. Complexe des accélérateurs du CERN - Août 2018", https://cds.cern.ch/record/2636343. General Photo. (Cited on page 35)

[71] H. Wiedemann, "Particle Accelerator Physics". Graduate Texts in Physics. Springer, Berlin, Germany, 2015, doi:10.1007/978-3-319-18317-6. ISBN: 978-3-319-18316-9, 978-3-319-18317-6. (Cited on page 34)

[72] O. S. Bruning et al., "LHC Design Report Vol.1: The LHC Main Ring", doi:10.5170/CERN-2004-003-V-1. (Cited on page 34)

[73] P. Lebrun, "Cryogenics for the Large Hadron Collider", *IEEE Transactions on Applied Superconductivity* **10** (2000), no. 1, doi:10.1109/77.828526. (Cited on page 34)

[74] ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider", *JINST* **3** (2008) S08003, doi:10.1088/1748-0221/3/08/S08003. (Cited on page 35)

[75] ALICE Collaboration, "The ALICE experiment at the CERN LHC", *JINST* **3** (2008) S08002, doi:10.1088/1748-0221/3/08/S08002. (Cited on page 35)

[76] CMS Collaboration, "The CMS Experiment at the CERN LHC", *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004. (Cited on pages 36, 41, 49, and 52)

[77] LHCb Collaboration, "The LHCb Detector at the LHC", *JINST* **3** (2008) S08005, doi:10.1088/1748-0221/3/08/S08005. (Cited on page 36)

[78] LHCf Collaboration, "The LHCf detector at the CERN Large Hadron Collider", *JINST* **3** (2008) S08006, doi:10.1088/1748-0221/3/08/S08006. (Cited on page 36)

[79] FASER Collaboration, "The FASER detector", *JINST* **19** (2024), no. 05, P05066, doi:10.1088/1748-0221/19/05/P05066, arXiv:2207.11427. (Cited on page 36)

[80] SND@LHC Collaboration, "SND@LHC: the scattering and neutrino detector at the LHC", *JINST* **19** (2024), no. 05, P05067, doi:10.1088/1748-0221/19/05/P05067, arXiv:2210.02784. (Cited on page 36)

[81] TOTEM Collaboration, "The TOTEM experiment at the CERN Large Hadron Collider", *JINST* **3** (2008) S08007, doi:10.1088/1748-0221/3/08/S08007. (Cited on page 36)

[82] J. L. Pinfold, "The MoEDAL Experiment — Searching for Highly Ionizing Particles at the LHC", in *New Trends in High-Energy Physics 2008*, p. 217. 2008. doi:10.1007/978-90-481-2287-5_21. (Cited on page 36)

[83] CERN, "HiLumi LHC: The High Luminosity LHC project", https://hilumilhc.web.cern.ch/. Accessed: Nov. 3, 2024. (Cited on page 38)

[84] CMS, "CMS Luminosity - Public Results", https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults/. Accessed: Nov. 4, 2024. (Cited on pages 39 and 40)

[85] CMS Collaboration, "Observation of the Higgs boson decay to a pair of τ leptons with the CMS detector", *Phys. Lett. B* **779** (2018) 283–316, doi:10.1016/j.physletb.2018.02.004, arXiv:1708.00373. (Cited on page 39)

[86] CMS Collaboration, "Observation of Higgs boson decay to bottom quarks", *Phys. Rev. Lett.* **121** (2018), no. 12, 121801, doi:10.1103/PhysRevLett.121.121801, arXiv:1808.08242. (Cited on page 39)

[87] G. Apollinari et al., "High-Luminosity Large Hadron Collider (HL-LHC) : Preliminary Design Report", doi:10.5170/CERN-2015-005. (Cited on page 40)

[88] ATLAS Collaboration, "Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment", Technical Report, LHCC-I-023, CERN-LHCC-2012-022, 2012. (Cited on page 40)

[89] CMS Collaboration, "Technical Proposal for the Phase-II Upgrade of the CMS Detector", Technical Report, CERN-LHCC-2015-010, LHCC-P-008, CMS-TDR-15-02, CERN, Geneva, 2015, doi:10.17181/CERN.VU8I.D59J. (Cited on pages 40, 57, and 71)

[90] LHCb Collaboration, "Framework TDR for the LHCb Upgrade II: Opportunities in flavour physics, and beyond, in the HL-LHC era", Technical Report, CERN-LHCC-2021-012, LHCB-TDR-023, 2021. (Cited on page 40)

[91] ALICE Collaboration, "Upgrade of the ALICE Experiment: Letter Of Intent", *J. Phys. G* **41** (2014) 087001, doi:10.1088/0954-3899/41/8/087001. (Cited on page 40)

[92] I. Neutelings, "CMS coordinate system", https://tikz.net/axis3d_cms/. Accessed: Nov. 5, 2024. (Cited on page 42)

[93] CMS Collaboration, "Development of the CMS detector for the CERN LHC Run 3", *JINST* **19** (2024), no. 05, P05064, doi:10.1088/1748-0221/19/05/P05064, arXiv:2309.05466. (Cited on pages 43, 46, 51, 52, 54, and 55)

[94] CMS Collaboration, "The CMS tracker system project: Technical Design Report", Technical Report, CERN-LHCC-98-006, CMS-TDR-5, CERN, 1997. (Cited on page 44)

[95] CMS Collaboration, "CMS Technical Design Report for the Pixel Detector Upgrade", Technical Report, CERN-LHCC-2012-016, CMS-TDR-011, CERN, 2012, doi:10.2172/1151650. (Cited on pages 44 and 45)

[96] CMS Tracker Group Collaboration, "The CMS Phase-1 Pixel Detector Upgrade", *JINST* **16** (2021), no. 02, P02027, doi:10.1088/1748-0221/16/02/P02027, arXiv:2012.14304. (Cited on pages 44, 45, and 46)

[97] CMS Collaboration, "The Phase-2 Upgrade of the CMS Tracker", Technical Report, CERN-LHCC-2017-009, CMS-TDR-014, CERN, Geneva, 2017, doi:10.17181/CERN.QZ28.FLHW. (Cited on pages 45 and 57)

[98] CMS Collaboration, "Pixel Detector Performance in Run 3 (June 2024)", CMS Detector Performance Summary, CMS-DP-2024-078, 2024. (Cited on page 46)

[99] CMS Collaboration, "CMS Silicon Strip Tracker Performance in 2023", CMS Detector Performance Summary, CMS-DP-2023-040, 2023. (Cited on page 47)

[100] CMS Collaboration, "Alignment of the CMS tracker with LHC and cosmic ray data", *JINST* **9** (2014) P06009, doi:10.1088/1748-0221/9/06/P06009, arXiv:1403.2286. (Cited on pages 47, 159, 160, 165, 166, and 167)

[101] CMS Collaboration, "Strategies and performance of the CMS silicon tracker alignment during LHC Run 2", *Nucl. Instrum. Meth. A* **1037** (2022) 166795, doi:10.1016/j.nima.2022.166795, arXiv:2111.08757. (Cited on pages 47, 159, 161, 165, and 166)

[102] CMS Collaboration, "The CMS electromagnetic calorimeter project: Technical Design Report", Technical Report, CERN-LHCC-97-033, CMS-TDR-4, 1997. (Cited on pages 48 and 49)

[103] P. Adzic et al., "Energy resolution of the barrel of the CMS electromagnetic calorimeter", *JINST* **2** (2007) P04004, doi:10.1088/1748-0221/2/04/P04004. (Cited on page 50)

[104] CMS Collaboration, "Performance of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 13$ TeV", *JINST* **19** (2024), no. 09, P09004, doi:10.1088/1748-0221/19/09/P09004, arXiv:2403.15518. (Cited on page 50)

[105] CMS Collaboration, "The CMS hadron calorimeter project: Technical Design Report", Technical Report, CERN-LHCC-97-031, CMS-TDR-2, 1997. (Cited on pages 50 and 51)

[106] CMS Collaboration, "CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter", Technical Report, CERN-LHCC-2012-015, CMS-TDR-010, 9, 2012, doi:10.2172/1151651. (Cited on pages 51 and 52)

[107] CMS HCAL Collaboration, "Design, performance, and calibration of CMS hadron-barrel calorimeter wedges", *Eur. Phys. J. C* **55** (2008), no. 1, 159–171, doi:10.1140/epjc/s10052-008-0573-y. (Cited on page 52)

[108] CMS Collaboration, "Calibration of the CMS hadron calorimeters using proton-proton collision data at $\sqrt{s} = 13$ TeV", *JINST* **15** (2020), no. 05, P05002, doi:10.1088/1748-0221/15/05/P05002, arXiv:1910.00079. (Cited on page 52)

[109] CMS Collaboration, "The CMS muon project: Technical Design Report", Technical Report, CERN-LHCC-97-032, CMS-TDR-3, 1997. (Cited on page 52)

[110] CMS Collaboration, "CMS Technical Design Report for the Muon Endcap GEM Upgrade", Technical Report, CERN-LHCC-2015-012, CMS-TDR-013, 2015. (Cited on page 53)

[111] CMS Collaboration, "Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV", *JINST* **13** (2018), no. 06, P06015, doi:10.1088/1748-0221/13/06/P06015, arXiv:1804.04528. (Cited on pages 53, 54, and 71)

[112] CMS Collaboration, "GEM performance results with 2024 data", CMS Detector Performance Summary, CMS-DP-2024-073, 2024. (Cited on page 55)

[113] CMS Collaboration, "CMS TriDAS project: Technical Design Report, Volume 1: The Trigger Systems", Technical Report, CERN-LHCC-2000-038, CMS-TDR-6-1, 2004. (Cited on page 55)

[114] CMS Collaboration, "Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JINST* **15** (2020), no. 10, P10017, doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165. (Cited on pages 55 and 56)

[115] CMS Collaboration, "Enriching the physics program of the CMS experiment via data scouting and data parking", Technical Report, , CERN, Geneva, 2024, arXiv:2403.16134, To be submitted to Physics Reports. (Cited on page 56)

[116] C. CMS, "A MIP Timing Detector for the CMS Phase-2 Upgrade", Technical Report, CERN-LHCC-2019-003, CMS-TDR-020, CERN, Geneva, 2019. (Cited on page 57)

[117] CMS Collaboration, "The Phase-2 Upgrade of the CMS Barrel Calorimeters", Technical Report, CERN-LHCC-2017-011, CMS-TDR-015, CERN, Geneva, 2017. (Cited on page 57)

[118] CMS Collaboration, "The Phase-2 Upgrade of the CMS Endcap Calorimeter", Technical Report, CERN-LHCC-2017-023, CMS-TDR-019, CERN, Geneva, 2017, doi:10.17181/CERN.IV8M.1JY2. (Cited on page 58)

[119] CMS Collaboration, "The Phase-2 Upgrade of the CMS Muon Detectors", Technical Report, CERN-LHCC-2017-012, CMS-TDR-016, 2017. (Cited on page 58)

[120] CMS Collaboration, "The Phase-2 Upgrade of the CMS Level-1 Trigger", Technical Report, CERN-LHCC-2020-004, CMS-TDR-021, CERN, Geneva, 2020. (Cited on page 58)

[121] C. Collaboration, "The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger", Technical Report, CERN-LHCC-2021-007, CMS-TDR-022, CERN, Geneva, 2021. (Cited on page 58)

[122] A. Buckley et al., "General-purpose event generators for LHC physics", *Phys. Rept.* **504** (2011) 145–233, doi:10.1016/j.physrep.2011.03.005, arXiv:1101.2599. (Cited on pages 60 and 64)

[123] CMS Collaboration, "Particle-flow reconstruction and global event description with the CMS detector", *JINST* **12** (2017), no. 10, P10003, doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965. (Cited on pages 61, 66, 67, and 75)

[124] J. C. Collins, D. E. Soper, and G. F. Sterman, "Factorization of Hard Processes in QCD", *Adv. Ser. Direct. High Energy Phys.* **5** (1989) 1–91, doi:10.1142/9789814503266_0001, arXiv:hep-ph/0409313. (Cited on page 60)

[125] NNPDF Collaboration, "Parton distributions from high-precision collider data", *Eur. Phys. J. C* **77** (2017), no. 10, 663, doi:10.1140/epjc/s10052-017-5199-5, arXiv:1706.00428. (Cited on pages 62, 63, and 88)

[126] V. V. Sudakov, "Vertex parts at very high-energies in quantum electrodynamics", *Sov. Phys. JETP* **3** (1956) 65–71. (Cited on page 62)

[127] B. Andersson, G. Gustafson, and B. Soderberg, "A General Model for Jet Fragmentation", *Z. Phys. C* **20** (1983) 317, doi:10.1007/BF01407824. (Cited on page 64)

[128] T. Sjostrand, "Jet Fragmentation of Nearby Partons", *Nucl. Phys. B* **248** (1984) 469–502, doi:10.1016/0550-3213(84)90607-2. (Cited on page 64)

[129] R. D. Field and S. Wolfram, "A QCD Model for e+ e- Annihilation", *Nucl. Phys. B* **213** (1983) 65–84, doi:10.1016/0550-3213(83)90175-X. (Cited on page 64)

[130] CMS Collaboration, "Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements", *Eur. Phys. J. C* **80** (2020), no. 1, 4, doi:10.1140/epjc/s10052-019-7499-4, arXiv:1903.12179. (Cited on pages 64 and 88)

[131] C. Bierlich et al., "A comprehensive guide to the physics and usage of PYTHIA 8.3", *SciPost Phys. Codeb.* **2022** (2022) 8, doi:10.21468/SciPostPhysCodeb.8, arXiv:2203.11601. (Cited on pages 64 and 86)

[132] T. Sjöstrand et al., "An introduction to PYTHIA 8.2", *Comput. Phys. Commun.* **191** (2015) 159–177, doi:10.1016/j.cpc.2015.01.024, arXiv:1410.3012. (Cited on pages 64 and 86)

[133] J. Alwall et al., "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations", *JHEP* **07** (2014) 079, doi:10.1007/JHEP07(2014)079, arXiv:1405.0301. (Cited on pages 65 and 86)

[134] J. Alwall et al., "Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions", *Eur. Phys. J. C* **53** (2008) 473–500, doi:10.1140/epjc/s10052-007-0490-5, arXiv:0706.2569. (Cited on pages 65 and 88)

[135] R. Frederix and S. Frixione, "Merging meets matching in MC@NLO", *JHEP* **12** (2012) 061, doi:10.1007/JHEP12(2012)061, arXiv:1209.6215. (Cited on pages 65 and 88)

[136] P. Nason, "A New method for combining NLO QCD with shower Monte Carlo algorithms", *JHEP* **11** (2004) 040, doi:10.1088/1126-6708/2004/11/040, arXiv:hep-ph/0409146. (Cited on pages 65 and 88)

[137] S. Frixione, P. Nason, and C. Oleari, "Matching NLO QCD computations with Parton Shower simulations: the POWHEG method", *JHEP* **11** (2007) 070, doi:10.1088/1126-6708/2007/11/070, arXiv:0709.2092. (Cited on pages 65 and 88)

[138] S. Alioli, P. Nason, C. Oleari, and E. Re, "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX", *JHEP* **06** (2010) 043, doi:10.1007/JHEP06(2010)043, arXiv:1002.2581. (Cited on pages 65 and 88)

[139] GEANT4 Collaboration, "GEANT4 - A Simulation Toolkit", *Nucl. Instrum. Meth. A* **506** (2003) 250–303, doi:10.1016/S0168-9002(03)01368-8. (Cited on page 65)

[140] J. Allison et al., "Geant4 developments and applications", *IEEE Trans. Nucl. Sci.* **53** (2006) 270, doi:10.1109/TNS.2006.869826. (Cited on page 65)

[141] CMS Collaboration, "The fast simulation of the CMS detector at LHC", *J. Phys. Conf. Ser.* **331** (2011) 032049, doi:10.1088/1742-6596/331/3/032049. (Cited on page 65)

[142] A. Giammanco, "The Fast Simulation of the CMS Experiment", *J. Phys. Conf. Ser.* **513** (2014) 022012, doi:10.1088/1742-6596/513/2/022012. (Cited on page 65)

[143] CMS Collaboration, "Recent Developments in CMS Fast Simulation", *PoS* **ICHEP2016** (2016) 181, doi:10.22323/1.282.0181, arXiv:1701.03850. (Cited on page 65)

[144] CMS Collaboration, "Refining fast simulation using machine learning", *EPJ Web Conf.* **295** (2024) 09032, doi:10.1051/epjconf/202429509032, arXiv:2309.12919. (Cited on page 65)

[145] CMS Collaboration, "Description and performance of track and primary-vertex reconstruction with the CMS tracker", *JINST* **9** (2014), no. 10, P10009,

doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569. (Cited on pages 68 and 70)

[146] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems", *J. Fluids Eng.* **82** (1960), no. 1, 35–45, doi:10.1115/1.3662552. (Cited on page 68)

[147] P. Billoir, "Progressive track recognition with a Kalman like fitting procedure", *Comput. Phys. Commun.* **57** (1989) 390–394, doi:10.1016/0010-4655(89)90249-X. (Cited on page 68)

[148] P. Billoir and S. Qian, "Simultaneous pattern recognition and track fitting by the Kalman filtering method", *Nucl. Instrum. Meth. A* **294** (1990) 219–228, doi:10.1016/0168-9002(90)91835-Y. (Cited on page 68)

[149] R. Mankel, "A Concurrent track evolution algorithm for pattern recognition in the HERA-B main tracking system", *Nucl. Instrum. Meth. A* **395** (1997) 169–184, doi:10.1016/S0168-9002(97)00705-5. (Cited on page 68)

[150] R. Fruhwirth, "Application of Kalman filtering to track and vertex fitting", *Nucl. Instrum. Meth. A* **262** (1987) 444–450, doi:10.1016/0168-9002(87)90887-4. (Cited on page 68)

[151] CMS Collaboration, "2017 tracking performance plots", CMS Detector Performance Summary, CMS-DP-2017-015, 2017. (Cited on page 69)

[152] CMS Collaboration, "Muon tracking performance in the CMS Run-2 Legacy data using the tag-and-probe technique", CMS Detector Performance Summary, CMS-DP-2020-035, 2020. (Cited on pages 69 and 70)

[153] CMS Collaboration, "Performance of Run 3 track reconstruction with the mkFit algorithm", CMS Detector Performance Summary, CMS-DP-2022-018, 2022. (Cited on page 69)

[154] CMS Collaboration, "Performance of the track selection DNN in Run 3", CMS Detector Performance Summary, CMS-DP-2023-009, 2023. (Cited on page 69)

[155] K. Rose, "Deterministic annealing for clustering, compression, classification, regression, and related optimization problems", *IEEE Proc.* **86** (1998), no. 11, 2210–2239, doi:10.1109/5.726788. (Cited on page 71)

[156] T. Speer et al., "Vertex Fitting in the CMS Tracker", Technical Report, CMS-NOTE-2006-032, CERN, Geneva, 2006. (Cited on page 71)

[157] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_{T} jet clustering algorithm", *JHEP* **04** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189. (Cited on pages 71 and 74)

[158] CMS Collaboration, "Performance of CMS Muon Reconstruction in pp Collision Events at $\sqrt{s} = 7$ TeV", *JINST* **7** (2012) P10002, doi:10.1088/1748-0221/7/10/P10002, arXiv:1206.4071. (Cited on page 71)

[159] CMS Collaboration, "Muon identification using multivariate techniques in the CMS experiment in proton-proton collisions at sqrt(s) = 13 TeV", *JINST* **19** (2024),

no. 02, P02031, doi:10.1088/1748-0221/19/02/P02031, arXiv:2310.03844. (Cited on page 73)

[160] CMS Collaboration, "Muon identification and isolation efficiencies with 2017 and 2018 data", CMS Detector Performance Summary, CMS-DP-2018-042, 2018. (Cited on page 73)

[161] CMS Collaboration, "Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC", *JINST* **16** (2021), no. 05, P05014, doi:10.1088/1748-0221/16/05/P05014, arXiv:2012.06888. (Cited on pages 73 and 74)

[162] CMS Collaboration, "Performance of Electron Reconstruction and Selection with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV", *JINST* **10** (2015), no. 06, P06005, doi:10.1088/1748-0221/10/06/P06005, arXiv:1502.02701. (Cited on page 74)

[163] W. Adam, R. Fruhwirth, A. Strandlie, and T. Todorov, "Reconstruction of electrons with the Gaussian sum filter in the CMS tracker at LHC", *eConf* **C0303241** (2003) TULT009, doi:10.1088/0954-3899/31/9/N01, arXiv:physics/0306087. (Cited on page 74)

[164] M. Cacciari, G. P. Salam, and G. Soyez, "FastJet user manual", *Eur. Phys. J. C* **72** (2012) 1896, doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097. (Cited on page 74)

[165] CMS Collaboration, "Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS", *JINST* **6** (2011) P11002, doi:10.1088/1748-0221/6/11/P11002, arXiv:1107.4277. (Cited on pages 74, 75, and 76)

[166] CMS Collaboration, "Jet algorithms performance in 13 TeV data", Technical Report, CMS-PAS-JME-16-003, CERN, Geneva, 2017. (Cited on pages 74 and 75)

[167] CMS Collaboration, "Pileup mitigation at CMS in 13 TeV data", *JINST* **15** (2020), no. 09, P09018, doi:10.1088/1748-0221/15/09/P09018, arXiv:2003.00503. (Cited on page 75)

[168] CMS Collaboration, "Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV", *JINST* **12** (2017), no. 02, P02014, doi:10.1088/1748-0221/12/02/P02014, arXiv:1607.03663. (Cited on page 76)

[169] CMS Collaboration, "Jet energy scale and resolution measurement with Run 2 Legacy Data Collected by CMS at 13 TeV", CMS Detector Performance Summary, CMS-DP-2021-033, 2021. (Cited on page 76)

[170] CMS Collaboration, "Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV", *JINST* **13** (2018), no. 05, P05011, doi:10.1088/1748-0221/13/05/P05011, arXiv:1712.07158. (Cited on page 77)

[171] E. Bols et al., "Jet Flavour Classification Using DeepJet", *JINST* **15** (2020), no. 12, P12012, doi:10.1088/1748-0221/15/12/P12012, arXiv:2008.10519. (Cited on page 77)

[172] CMS Collaboration, "Performance of the DeepJet b tagging algorithm using 41.9/fb of data from proton-proton collisions at 13TeV with Phase 1 CMS detector", CMS Detector Performance Summary, CMS-DP-2018-058, 2018. (Cited on page 77)

[173] CMS Collaboration, "Performance summary of AK4 jet b tagging with data from proton-proton collisions at 13 TeV with the CMS detector", CMS Detector Performance Summary, CMS-DP-2023-005, 2023. (Cited on pages 77 and 123)

[174] CMS Collaboration, "Performance of tau-lepton reconstruction and identification in CMS", *JINST* **7** (2012) P01001, doi:10.1088/1748-0221/7/01/P01001, arXiv:1109.6034. (Cited on page 78)

[175] CMS Collaboration, "Reconstruction and identification of τ lepton decays to hadrons and ν_τ at CMS", *JINST* **11** (2016), no. 01, P01019, doi:10.1088/1748-0221/11/01/P01019, arXiv:1510.07488. (Cited on page 78)

[176] CMS Collaboration, "Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV", *JINST* **13** (2018), no. 10, P10005, doi:10.1088/1748-0221/13/10/P10005, arXiv:1809.02816. (Cited on pages 78 and 79)

[177] CMS Collaboration, "Identification of hadronic tau lepton decays using a deep neural network", *JINST* **17** (2022) P07023, doi:10.1088/1748-0221/17/07/P07023, arXiv:2201.08458. (Cited on page 79)

[178] CMS Collaboration, "Performance of the CMS missing transverse momentum reconstruction in pp data at $\sqrt{s} = 8$ TeV", *JINST* **10** (2015), no. 02, P02006, doi:10.1088/1748-0221/10/02/P02006, arXiv:1411.0511. (Cited on page 80)

[179] CMS Collaboration, "Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector", *JINST* **14** (2019), no. 07, P07004, doi:10.1088/1748-0221/14/07/P07004, arXiv:1903.06078. (Cited on page 80)

[180] ATLAS Collaboration, "Search for new light gauge bosons in Higgs boson decays to four-lepton final states in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector at the LHC", *Phys. Rev. D* **92** (2015) 092001, doi:10.1103/PhysRevD.92.092001, arXiv:1505.07645. (Cited on page 82)

[181] CMS Collaboration, "Search for a Non-Standard-Model Higgs Boson Decaying to a Pair of New Light Bosons in Four-Muon Final States", *Phys. Lett. B* **726** (2013) 564, doi:10.1016/j.physletb.2013.09.009, arXiv:1210.7619. (Cited on page 82)

[182] CMS Collaboration, "A search for pair production of new light bosons decaying into muons", *Phys. Lett. B* **752** (2016) 146, doi:10.1016/j.physletb.2015.10.067, arXiv:1506.00424. (Cited on page 82)

[183] CMS Collaboration, "Search for light bosons in decays of the 125 GeV Higgs boson in proton-proton collisions at $\sqrt{s} = 8$ TeV", *JHEP* **10** (2017) 076, doi:10.1007/JHEP10(2017)076, arXiv:1701.02032. (Cited on page 82)

[184] CMS Collaboration, "Search for a very light NMSSM Higgs boson produced in decays of the 125 GeV scalar boson and decaying into τ leptons in pp collisions at $\sqrt{s} = 8$ TeV", *JHEP* **01** (2016) 079, doi:10.1007/JHEP01(2016)079, arXiv:1510.06534. (Cited on page 82)

[185] ATLAS Collaboration, "Search for new phenomena in events with at least three photons collected in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector", *Eur. Phys. J. C* **76** (2016) 210, doi:10.1140/epjc/s10052-016-4034-8, arXiv:1509.05051. (Cited on page 82)

[186] ATLAS Collaboration, "Search for Higgs boson decays to beyond-the-standard-model light bosons in four-lepton events with the ATLAS detector at $\sqrt{s}$ =13 TeV", *JHEP* **06** (2018) 166, doi:10.1007/JHEP06(2018)166, arXiv:1802.03388. (Cited on page 82)

[187] CMS Collaboration, "A search for pair production of new light bosons decaying into muons in proton-proton collisions at 13 TeV", *Phys. Lett. B* **796** (2019) 131, doi:10.1016/j.physletb.2019.07.013, arXiv:1812.00380. (Cited on page 82)

[188] ATLAS Collaboration, "Search for Higgs bosons decaying to aa in the $\mu\mu\tau\tau$ final state in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS experiment", *Phys. Rev. D* **92** (2015) 052002, doi:10.1103/PhysRevD.92.052002, arXiv:1505.01609. (Cited on page 82)

[189] CMS Collaboration, "Search for an exotic decay of the Higgs boson to a pair of light pseudoscalars in the final state of two muons and two τ leptons in proton-proton collisions at $\sqrt{s}$ =13 TeV", *JHEP* **11** (2018) 018, doi:10.1007/JHEP11(2018)018, arXiv:1805.04865. (Cited on page 82)

[190] CMS Collaboration, "Search for a light pseudoscalar Higgs boson in the boosted $\mu\mu\tau\tau$ final state in proton-proton collisions at $\sqrt{s}$ =13 TeV", *JHEP* **08** (2020) 139, doi:10.1007/JHEP08(2020)139, arXiv:2005.08694. (Cited on page 82)

[191] ATLAS Collaboration, "Search for Higgs boson decays into a pair of light bosons in the $bb\mu\mu$ final state in pp collision at $\sqrt{s}$ =13 TeV with the ATLAS detector", *Phys. Lett. B* **790** (2019) 1, doi:10.1016/j.physletb.2018.10.073, arXiv:1807.00539. (Cited on page 82)

[192] CMS Collaboration, "Search for exotic decays of the Higgs boson to a pair of pseudoscalars in the $\mu\mu$bb and $\tau\tau$bb final states", *Eur. Phys. J. C* **84** (2024) 493, doi:10.1140/epjc/s10052-024-12727-4, arXiv:2402.13358. (Cited on page 82)

[193] ATLAS Collaboration, "Search for Higgs boson decays into pairs of light (pseudo)scalar particles in the $\gamma\gamma jj$ final state in pp collisions at $\sqrt{s}$ =13 TeV with the ATLAS detector", *Phys. Lett. B* **782** (2018) 750, doi:10.1016/j.physletb.2018.06.011, arXiv:1803.11145. (Cited on page 82)

[194] CMS Collaboration, "Search for the decay of the Higgs boson to a pair of light pseudoscalar bosons in the final state with four bottom quarks in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **06** (2024) 097, doi:10.1007/JHEP06(2024)097, arXiv:2403.10341. (Cited on page 82)

[195] CMS Collaboration, "Search for the exotic decay of the Higgs boson into two light pseudoscalars with four photons in the final state in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **07** (2023) 148, doi:10.1007/JHEP07(2023)148, arXiv:2208.01469. (Cited on page 82)

[196] CMS Collaboration, "Search for exotic Higgs boson decays $H \to \mathcal{A}\mathcal{A} \to 4\gamma$ with events containing two merged diphotons in proton-proton collisions at $\sqrt{s} = 13$ TeV", *Phys. Rev. Lett.* **131** (2023) 101801, doi:10.1103/PhysRevLett.131.101801, arXiv:2209.06197. (Cited on page 82)

[197] CMS Collaboration, "Search for light pseudoscalar boson pairs produced from decays of the 125 GeV Higgs boson in final states with two muons and two nearby tracks in pp collisions at $\sqrt{s} = 13$ TeV", *Phys. Lett. B* **800** (2020) 135087, doi:10.1016/j.physletb.2019.135087, arXiv:1907.07235. (Cited on pages 83, 144, and 156)

[198] CMS Collaboration, "Noise Filter Recommendations for Run II", https://twiki.cern.ch/twiki/bin/viewauth/CMS/MissingETOptionalFiltersRun2. Accessed: Jan. 28, 2025. (Cited on page 92)

[199] CMS Collaboration, "HEM Issue in 2018", https://twiki.cern.ch/twiki/bin/view/CMS/SUSRecommendationsRun2UltraLegacy#HEM_issue_in_2018. Accessed: Jan. 28, 2025. (Cited on page 92)

[200] M. Lisanti and J. G. Wacker, "Discovering the Higgs boson with low mass muon pairs", *Phys. Rev. D* **79** (2010) 115006, doi:10.1103/PhysRevD.79.115006, arXiv:0903.1377. (Cited on page 99)

[201] CMS Collaboration, "L1 Prefiring Weights Recipe", https://twiki.cern.ch/twiki/bin/viewauth/CMS/L1PrefiringWeightRecipe. Accessed: Jan. 28, 2025. (Cited on page 100)

[202] CMS Collaboration, "Rochester Corrections", https://twiki.cern.ch/twiki/bin/view/CMS/RochcorMuon. Accessed: Jan. 28, 2025. (Cited on page 100)

[203] G. Bozzi, S. Catani, D. de Florian, and M. Grazzini, "Transverse-momentum resummation and the spectrum of the Higgs boson at the LHC", *Nucl. Phys. B* **737** (2006) 73, doi:10.1016/j.nuclphysb.2005.12.022, arXiv:hep-ph/0508068. (Cited on page 106)

[204] D. de Florian, G. Ferrera, M. Grazzini, and D. Tommasini, "Transverse-momentum resummation: Higgs boson production at the Tevatron and the LHC", *JHEP* **11** (2011) 064, doi:10.1007/JHEP11(2011)064, arXiv:1109.2109. (Cited on page 106)

[205] CMS Collaboration, "Recommendations for calculating b-tagging scale factors", https://btv-wiki.docs.cern.ch/PerformanceCalibration/fixedWPSFRecommendations/. Accessed: Jan. 29, 2025. (Cited on page 106)

[206] R. J. Barlow and C. Beeston, "Fitting using finite Monte Carlo samples", *Comput. Phys. Commun.* **77** (1993) 219–228, doi:10.1016/0010-4655(93)90005-W. (Cited on pages 122 and 127)

[207] NNPDF Collaboration, "Parton distributions for the LHC Run II", *JHEP* **04** (2015) 040, doi:10.1007/JHEP04(2015)040, arXiv:1410.8849. (Cited on page 124)

[208] J. Pumplin et al., "New generation of parton distributions with uncertainties from global QCD analysis", *JHEP* **07** (2002) 012, doi:10.1088/1126-6708/2002/07/012, arXiv:hep-ph/0201195. (Cited on page 124)

[209] J. S. Conway, "Incorporating Nuisance Parameters in Likelihoods for Multisource Spectra", in *PHYSTAT 2011*, pp. 115–120. 2011. arXiv:1103.0354. doi:10.5170/CERN-2011-006.115. (Cited on page 127)

[210] J. Neyman and E. S. Pearson, "On the Problem of the Most Efficient Tests of Statistical Hypotheses", *Phil. Trans. Roy. Soc. Lond. A* **231** (1933), no. 694-706, 289–337, doi:10.1098/rsta.1933.0009. (Cited on page 128)

[211] A. L. Read, "Presentation of search results: The CL_s technique", *J. Phys. G* **28** (2002) 2693–2704, doi:10.1088/0954-3899/28/10/313. (Cited on page 129)

[212] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics", *Eur. Phys. J. C* **71** (2011) 1554, doi:10.1140/epjc/s10052-011-1554-0, arXiv:1007.1727. [Erratum: Eur.Phys.J.C 73, 2501 (2013)]. (Cited on page 130)

[213] S. S. Wilks, "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses", *Annals Math. Statist.* **9** (1938), no. 1, 60–62, doi:10.1214/aoms/1177732360. (Cited on page 130)

[214] CMS Collaboration, "The CMS Statistical Analysis and Combination Tool: Combine", *Comput. Softw. Big Sci.* **8** (2024), no. 1, 19, doi:10.1007/s41781-024-00121-4, arXiv:2404.06614. (Cited on page 130)

[215] U. Haisch, J. F. Kamenik, A. Malinauskas, and M. Spira, "Collider constraints on light pseudoscalars", *JHEP* **03** (2018) 178, doi:10.1007/JHEP03(2018)178, arXiv:1802.02156. (Cited on pages 144 and 146)

[216] LHC Higgs Working Group Collaboration, "Scans in 2HDM+S parameter space", https://twiki.cern.ch/twiki/bin/view/LHCPhysics/LHCHWGExoticDecay?redirectedfrom=LHCPhysics.LHCHXSWGExoticDecay. Accessed: Jan. 28, 2025. (Cited on page 144)

[217] M. Baumgart and A. Katz, "Implications of a new light scalar near the bottomonium regime", *JHEP* **08** (2012) 133, doi:10.1007/JHEP08(2012)133, arXiv:1204.6032. (Cited on page 146)

[218] N. Bartosik, "Associated Top-Quark-Pair and b-Jet Production in the Dilepton Channel at $\sqrt{s} = 8$ TeV as Test of QCD and Background to tt+Higgs Production". PhD thesis, Hamburg U., 2015, https://cds.cern.ch/record/2047049. Presented 03 Jul 2015. (Cited on pages 161, 163, and 164)

[219] V. Blobel and C. Kleinwort, "A New method for the high precision alignment of track detectors", in *Conference on Advanced Statistical Techniques in Particle Physics*. 6, 2002. arXiv:hep-ex/0208021. (Cited on page 161)

[220] C. Kleinwort, "General Broken Lines as advanced track fitting method", *Nucl. Instrum. Meth. A* **673** (2012) 107–110, doi:10.1016/j.nima.2012.01.024, arXiv:1201.4320. (Cited on page 162)

[221] C. Kleinwort, "General Broken Lines Website", https://www.terascale.de/wiki/generalbrokenlines. Accessed: Jan. 28, 2025. (Cited on page 162)

[222] S.-C. T. Choi, C. C. Paige, and M. A. Saunders, "MINRES-QLP: A Krylov Subspace Method for Indefinite or Singular Symmetric Systems", *SIAM Journal on Scientific Computing* **33** (2011), no. 4, 1810–1836, doi:10.1137/100787921, arXiv:https://doi.org/10.1137/100787921. (Cited on page 162)

[223] S.-C. T. Choi and M. A. Saunders, "Algorithm 937: MINRES-QLP for symmetric and Hermitian linear equations and least-squares problems", *ACM Trans. Math. Softw.* **40** (March, 2014) doi:10.1145/2527267. (Cited on page 162)

[224] E. Anderson et al., "LAPACK Users' Guide". Society for Industrial and Applied Mathematics, third edition, 1999, doi:10.1137/1.9780898719604. (Cited on page 162)

[225] CMS Collaboration, "Alignment of the CMS Silicon Strip Tracker during stand-alone Commissioning", *JINST* **4** (2009) T07001, doi:10.1088/1748-0221/4/07/T07001, arXiv:0904.1220. (Cited on page 165)

[226] J. Draeger, "Track based alignment of the CMS silicon tracker and its implication on physics performance". PhD thesis, Hamburg U., 2011. doi:10.3204/DESY-THESIS-2011-026. (Cited on page 165)

[227] CMS Collaboration, "First Tracker Alignment results with 2021 cosmic ray data", CMS Detector Performance Summary, CMS-DP-2021-025, 2021. (Cited on page 168)

[228] CMS Collaboration, "Performance of CMS Tracker Alignment during LHC Run 2 and Commissioning towards Run 3", *J. Phys. Conf. Ser.* **2438** (2023), no. 1, 012112, doi:10.1088/1742-6596/2438/1/012112. (Cited on page 177)

[229] CMS Collaboration, "Automated alignment calibration in CMS during Run 3", *PoS* **EPS-HEP2023** (2024) 595, doi:10.22323/1.449.0595. (Cited on page 185)

Acknowledgments

This thesis is the culmination of research carried out over three and a half years. During this time, I have been fortunate enough to work alongside and learn from many exceptional individuals whose guidance, support, and encouragement have been instrumental in bringing this work to completion. I would like to take this opportunity to express my appreciation to all those who have contributed to this journey, both academically and personally.

The deepest gratitude goes to my supervisor, Dr. Alexei Raspereza, for his unwavering guidance and support throughout this period. Under his mentorship, I was given the freedom to develop independence in my work while always knowing that his advice and expertise were available whenever I needed them. His patience and dedication were invaluable, and he always took the time to provide clarity, no matter how small or trivial my questions were. Every conversation with him was a learning experience, and his passion for physics has been a constant source of inspiration. It was a privilege to be his student, and I could not have asked for a better supervisor.

I would also like to thank my co-supervisor, Prof. Dr. Elisabetta Gallo, who, despite her extensive responsibilities at DESY and within CMS, remained consistently supportive. Her input was critical at every stage and helped refine this work. She provided insight not only into the analysis but also encouraged me to participate in external events that widened my perspective beyond my immediate research.

A special mention goes to Dr. Sandra Consuegra Rodriguez for her continuous involvement throughout my PhD. As a newcomer in CMS, I relied heavily on her guidance not only to get started with the analysis but also to understand the complexities of the alignment of the CMS tracker. She was always available to answer my questions and never hesitated to take the time to help. I am also grateful to her for entrusting me with responsibilities within the tracker alignment group and, later, the Alignment and Calibration Database group.

I am especially thankful to Dr. Henriette Aarup Petersen. I really enjoyed working with her on tracker alignment. Though the work was often challenging, our conversations about the tracker were always insightful, and I learned a great deal. She has consistently pushed for my growth within the tracker group, and I'm genuinely grateful for that.

Dr. Andrea Cardini has been a pillar of support throughout my PhD. Even though he was not directly involved in my work, he was always available to offer advice and guidance on multiple aspects of my professional life. A special thank you also goes to Dr. Danyer Pérez Adán for his help with the analysis and for always taking the time to discuss ideas and clarify doubts.

I extend a heartfelt thank you to Dr. Rainer Mankel and Dr. Roberval Walsh for their insightful feedback on my analysis during the numerous Higgs working meetings. Rainer's expertise in tracker alignment has also helped me gain a deeper understanding of the topic. I would further like to thank Dr. Olaf Behnke for the many lovely chats we've had over the years and for his invaluable help with the German abstract of this thesis.

I am really grateful to Mykyta for being an amazing office mate and friend throughout these years. We shared our entire PhD journey with each other, and it was always nice to have someone to talk to about both the highs and the lows along the way. I'm also grateful to my "newer" office mates, Saverio and Soumyaa, for bringing such a warm and enjoyable atmosphere to the office. Thank you for all the laughs and conversations.

While it is difficult to list everyone by name due to space constraints, I would also like to extend my sincere appreciation to everyone at DESY, be it the staff members, postdocs, or fellow PhD students, for fostering a collaborative, intellectually stimulating, and welcoming environment throughout the course of my doctoral studies.

Finally, I would like to say a huge thank you to my parents, who have always supported every decision that I've made, including moving to a faraway country to pursue my PhD. Thank you for inspiring me to follow my passion. I would also like to thank my closest friends, Arya and Lithin, who remain a constant presence in my life. They were always there to talk, through the good and the bad, share a laugh, and make it easier to wind down at the end of the day. Also, thanks to my long-time friends Krishna and Neha, whose presence in Europe and occasional meetups made the distance from home feel less stark. Lastly, I would like to extend my thanks to Vaisakh. Words cannot describe how much your constant support and presence mean to me. Thank you for making life a better place.