



FACULTÉ
DES SCIENCES



UNIVERSITÉ LIBRE DE BRUXELLES

**Search for new high mass resonances
or quantum black holes decaying to
lepton flavor violating final states
with the CMS detector**

Thesis presented by Diego BEGHIN

in fulfilment of the requirements of the PhD Degree in physical science
("Docteur en sciences physiques")

Année académique 2020-2021

Supervisor : Professor Barbara CLERBAUX
[IIHE]

Thesis jury :

Pascal VANLAER (Université libre de Bruxelles, Chair)
Laurent THOMAS (Université libre de Bruxelles, Secretary)
Barbara CLERBAUX (Université libre de Bruxelles)
Christophe DELAERE (Université catholique de Louvain)
Thomas HAMBYE (Université libre de Bruxelles)
Maria CEPEDA HERMIDA (Madrid-CIEMAT)



Abstract

This thesis presents a novel search for new physics in lepton flavor violating final states, using the CMS detector. There is first a discussion of the Standard Model of particle physics, with a particular emphasis on the issue of lepton flavor conservation, and how often that conservation is violated when generic modifications of the Standard Model are performed. The questions left unanswered by the Standard Model are presented, as well as some new physics models which resolve them, and in so doing imply the existence of processes violating charged lepton flavor conservation. R-parity violating supersymmetry, models with large extra dimensions, and models with an extra U(1) gauge symmetry are discussed. The Large Hadron Collider, the CMS detector, its trigger system and event reconstruction software are all described in detail. The data used for the analysis corresponds to an integrated luminosity of 137 fb^{-1} collected by CMS during three years of LHC Run 2 (2016-18), at a center-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. The background estimation strategy is to use simulations for all processes with real leptons, while large-yield processes with misidentified leptons are estimated with data-driven methods. The simulations are corrected by scale factors accounting for experimental calibration and more precise theoretical calculations. After comparing the observed data and the expected background, no evidence of new physics is found, and a statistical analysis is performed to exclude new physics models at the 95 % confidence level. For the benchmark lepton flavor violating Z' model, the lower limit in the $e\mu$ (respectively $e\tau$, $\mu\tau$) final state on the resonance mass is 5.0 TeV (resp. 4.2 TeV, 4.1 TeV). Such an analysis had never been performed in the tau channels with the CMS detector, and in all channels these results considerably improve on the previous state-of-the-art results.

Résumé

Cette thèse présente une analyse de données collectées par le détecteur CMS, dont le but est la recherche de nouvelle physique dans des états finaux qui violent la conservation de la saveur leptonique. Elle commence avec une discussion du Modèle Standard de la physique des particules, en insistant sur la conservation de la saveur leptonique dans ce modèle, et sur la violation de cette conservation qui découle d'un grand nombre de modifications du Modèle Standard. Les problèmes non résolus en physique des particules sont présentés, ainsi que certains des modèles de nouvelle physique qui y répondent, et qui impliquent une violation de la conservation de la saveur leptonique. Nous discutons notamment la supersymétrie avec violation de la R-parité, des modèles avec des dimensions supplémentaires, et des modèles avec une symétrie de jauge $U(1)$ supplémentaire. Ensuite, le LHC, le détecteur CMS, son système de déclenchement et son logiciel de reconstruction des événements sont tous décrits en détail. Les données analysées correspondent à une luminosité intégrée de 137 fb^{-1} collectée par CMS pendant trois années du Run 2 du LHC (2016-18), à une énergie du centre de masse $\sqrt{s} = 13\text{ TeV}$. La stratégie d'estimation du bruit de fond est de simuler tous les processus avec de vrais leptons dans l'état final, mais des méthodes basées sur les données sont utilisées pour les principaux bruits de fond où les leptons sont en fait des particules mal identifiées. Des facteurs de correction sont appliqués aux simulations afin de tenir compte d'effets expérimentaux et de calculs théoriques plus précis. Les données s'avèrent être compatibles avec les estimations du bruit de fond, en tenant compte des incertitudes, et une analyse statistique est faite afin d'exclure des modèles de nouvelle physique à un niveau de confiance de 95 %. Pour le modèle de référence, un boson Z' avec des désintégrations qui ne conservent pas la saveur leptonique, les limites inférieures sur la masse de la résonance sont respectivement 5.0 TeV, 4.2 TeV et 4.1 TeV pour les états finaux $e\mu$, $e\tau$ et $\mu\tau$. C'est la première analyse CMS de ce genre dans les états finaux $e\tau$ et $\mu\tau$, et dans tous les canaux les résultats sont des améliorations considérables par rapport aux meilleurs résultats précédents.

Contents

Introduction	ix
I Theory overview and experimental state of the art	1
1 The Standard Model of particle physics and lepton flavor conservation	3
1.1 Why the Standard Model?	3
1.2 A more detailed view of the Standard Model	4
1.2.1 Standard Model particles and forces	4
1.2.2 Spontaneous symmetry breaking via the Brout-Englert-Higgs mechanism	7
1.2.3 The full Lagrangian of the Standard Model	9
1.2.4 Why lepton flavor violating processes are forbidden in the Standard Model	10
1.3 Standard Model predictions and experimental tests	12
1.4 Standard Model incompleteness and hints of new physics	14
2 Beyond the Standard Model, towards charged lepton flavor violation	17
2.1 R-parity violating supersymmetry	17
2.1.1 A short overview of supersymmetry theory	17
2.1.2 R-parity conservation and violation	18
2.2 Quantum black holes	20
2.2.1 Extra dimensions and the weakness of gravity	20
2.2.2 From extra dimensions to quantum black holes	21
2.3 Lepton flavor violating heavy H bosons	22
2.4 Lepton flavor violating heavy Z' bosons	23
2.5 The experimental state of the art in charged lepton flavor violation searches . . .	25
2.5.1 Low energy searches for lepton flavor violation	25
2.5.2 Lepton universality	26
2.5.3 High energy searches for lepton flavor violation	29
II Experimental hardware and software	31
3 The Large Hadron Collider and the Compact Muon Solenoid	33
3.1 The Large Hadron Collider	33
3.1.1 Physical layout of the LHC	33
3.1.2 Particle beams in the LHC	35

3.1.3	The LHC injection chain	36
3.1.4	Past and future operations timeline	37
3.2	The Compact Muon Solenoid	39
3.2.1	Overview of CMS and its coordinate system	39
3.2.2	The silicon tracker	40
3.2.3	The electromagnetic calorimeter	42
3.2.4	The hadronic calorimeter	43
3.2.5	The muon system	45
4	The CMS trigger system	51
4.1	Level-1 trigger	51
4.1.1	Calorimeter trigger	52
4.1.2	Muon trigger	53
4.1.3	Global trigger	55
4.2	High Level Trigger	56
4.2.1	HLT algorithms	56
4.2.2	HLT rates monitoring	57
4.2.3	Dataset overlap study	60
5	Object reconstruction in CMS	67
5.1	Reconstruction of tracks, vertices and calorimeter clusters	67
5.2	Muons	70
5.3	Electrons	72
5.4	Photons and hadrons	76
5.5	Jets	76
5.6	Missing transverse energy	78
5.7	Taus	79
5.7.1	Reconstruction of hadronic taus	79
5.7.2	Identification of hadronic taus	81
5.7.3	Study of the efficiency of the anti-jet discriminators	84
III	Analysis of lepton flavor violating final states	93
6	Datasets and simulations	95
6.1	Datasets	95
6.2	Simulations	96
6.2.1	General overview of event simulation for collider experiments	96
6.2.2	Simulation of lepton flavor violating new physics	97
6.2.3	Simulation of Standard Model processes	99
6.2.4	Normalization, pileup reweighting and other scale factors	100
7	Analysis strategy	103
7.1	Trigger paths and event selection	103
7.2	Discriminating variables: invariant mass and collinear mass	106
7.3	Background estimation	114
7.4	Tau identification: comparison between boosted decision trees and DeepTau algorithms	117

8	Data-driven background estimations	121
8.1	Data-driven background estimations in the $e\mu$ final state	121
8.1.1	Electron and muon fake rates	121
8.1.2	Validation with the same sign method	122
8.2	Backgrounds with jet $\rightarrow \tau$ fakes	124
8.2.1	The fake factor method	126
8.2.2	Fake factors from different control regions	127
8.2.3	Fake factor dependence on the light lepton p_T	128
8.2.4	Validation using simulated events	131
8.2.5	Fake factor dependence on jet flavor	136
9	Signal region events and statistical analysis	139
9.1	Systematic uncertainties	139
9.2	Lepton kinematics and mass plots	142
9.3	Statistical analysis and limits on new physics models	145
9.4	Discussion of the results and prospects	153
	Conclusion	159
A	Special relativity and particle physics conventions	161
B	Validation of the $t\bar{t}$ background estimation	165
C	Comparison of jet $\rightarrow \tau$ fakes in the $e\tau$ and $\mu\tau$ channels	179
	Bibliography	181

Introduction

Modern scientific instruments allow us to observe the universe at vastly different distance scales. Telescopes have been used to observe other planets in the solar system for centuries, and more recently deep space telescopes give us insight into galaxies and galaxy clusters, the largest structures in the universe. It is also possible to probe the smallest distance scales: optical microscopes reveal the inner machinery of cells, and electronic microscopes allow even individual molecules and atoms to be observed. Cosmic ray experiments and particle colliders allow us to probe matter even deeper, at the subatomic level. This is the realm of particle physics, where the smallest known constituents of matter are studied.

Particle physics is also known as high energy physics, due to the ever higher energy densities needed to probe the smallest distance scales. Over the last century, particle colliders have thus been made ever bigger and more powerful in order to achieve the highest center-of-mass energies at the point of collision. The machine that is currently at the world's energy frontier is the CERN Large Hadron Collider (LHC), located at the French-Swiss border near Geneva. It collided protons at an unprecedented 13 TeV center-of-mass energy from 2015 to 2018, and is expected to reach 14 TeV at some point in the near future. This thesis analyzes data collected by the CMS detector at one of the collision points of the LHC, and addresses some of the open questions in particle physics.

Particle physics is described by the Standard Model (SM), which was constructed progressively in the 1960s and 70s as a large number of experimental discoveries demanded new theoretical explanations. It provides the current best explanation for all phenomena relating to three of the four fundamental forces: electromagnetism, and the weak and strong nuclear forces. Electromagnetism and the weak force are described at high energy by the electroweak theory, while the strong force is described by quantum chromodynamics. Together, these two descriptions form the Standard Model, which has been very successful at explaining novel experimental phenomena. But there are still some lingering issues: most notably, experiments have found that neutrinos have non-zero masses, in contradiction with their SM masslessness. Neutrino masses can be accommodated in the SM framework, but there is still disagreement about how exactly it should be done, and they could instead be a sign of the existence of new physics. Other problems in the SM include its incompatibility with the fourth fundamental force, gravity, and the lack of a good description of the dark matter and dark energy observed at galactic and cosmological scales. A more contested issue is the hierarchy problem, which is related to the large difference between the electroweak and quantum gravity scales, and seems to impose some fine-tuning of the scalar Brout-Englert-Higgs boson mass parameter. There is some debate about whether this apparent fine-tuning is a problem, but in any case many new theories have been advanced which lack this feature.

There are three families of leptons in the SM: the negatively charged electrons, muons and

taus, and their respective chargeless neutrinos (electron, muon and tau neutrinos). All of these particles have their own anti-particles. An accidental feature of the SM with massless neutrinos is that lepton flavor is conserved, i.e. there is a conservation of the total number of electrons and electron neutrinos minus the total number of anti-electrons and electron anti-neutrinos (the same is true replacing electrons by muons or taus). This conservation is accidental since it is not explicitly imposed as a symmetry, it is simply a consequence of the specific way in which the SM is constructed. Small tweaks to the SM, such as non-zero neutrino masses, often result in lepton flavor non-conservation. In general, unless lepton flavor conservation is explicitly imposed as a symmetry, SM extensions generically feature lepton flavor violating phenomena. Currently, such violations have only been observed in the neutrino sector, but in many new physics models observable effects also exist in the charged lepton sector. The search for lepton flavor violating phenomena in colliders such as the LHC is a necessary complement to the many searches already performed in lower energy experiments probing e.g. muon decay.

The thesis is organized in three parts, with part I providing an overview of the theory and the experimental state of the art in lepton flavor violation searches, part II describing the machines, detectors and software necessary to carry out the experiments and collect the data, and part III summarizing the analysis of the three lepton flavor violating final states. Part I is divided into chapter 1 explaining the Standard Model and chapter 2 introducing some new physics models and summarizing the experimental state of the art. Part II contains chapter 3, which describes the LHC and the CMS detector, chapter 4, which describes the system used to trigger data acquisition in CMS, and chapter 5, which introduces the various algorithms used to reconstruct collision events in CMS, with a focus on those most relevant to the analysis. Part III starts with chapter 6, which lists the specific datasets and simulation samples used in the analysis, as well as providing some details into how the simulations are done, then comes chapter 7, which describes the analysis strategy in signal event selection and background estimation, followed by chapter 8, describing the data-driven background estimation techniques used for events with wrongly-identified leptons. The final chapter (chapter 9) shows comparison plots between observed and expected events in the signal-optimized data region, explains the statistical analysis performed, shows the results of the analysis, and discusses how they relate to other experiments.

Part I

Theory overview and experimental state of the art

The Standard Model of particle physics and lepton flavor conservation

Particle physics, also known as high energy physics, is the field that studies the most fundamental constituents of matter and their interactions. The high energy density conditions necessary to probe such phenomena can be recreated in particle colliders or accelerators, or observed in radioactive decays or cosmic ray experiments. Two new forces (the weak and strong nuclear forces) and an ever-growing number of particles were discovered by such experiments during the 20th and 21st centuries, culminating in the discovery of the H boson in 2012 at the LHC. All of these new particles fit into a theoretical framework called the Standard Model (SM). The SM explains an enormous variety of observations with great accuracy, but a certain number of open questions do remain.

This chapter starts in section 1.1 with a historical overview of the experimental discoveries and theoretical breakthroughs which motivated the construction of the SM. Section 1.2 provides a general description of the SM, with a subsection focusing on lepton flavor conservation. Section 1.3 discusses how theoretical predictions are made from the SM Lagrangian, and shows some examples of experimentally confirmed predictions. The chapter closes in section 1.4 with examples of the shortcomings of the SM and why there are searches for new physics.

1.1 Why the Standard Model?

The first subatomic particle to be discovered was the electron, in 1897, by J. J. Thomson [1]. He showed that what were then known as cathode rays were composed of corpuscles much lighter than atoms, and that they were negatively charged¹. Many other discoveries were made in the following decades. By the 1950s, protons, neutrons, neutrinos, anti-electrons, muons, pions, and kaons were discovered, among others. Some were predicted by theory, some were unexpected. New models were needed in order to explain what was then called the “particle zoo”.

Regularities were observed in a subset of these particles that we now call hadrons. They could be classified according to their quantum numbers, like isospin² and strangeness [2, 3]. Those properties could be explained if hadrons were not fundamental particles, and were instead made up of smaller constituents called quarks. Evidence from deep inelastic scattering experiments, where high-energy electrons are made to collide into protons, showed that protons do have internal structure in the form of pointlike particles [4], which were identified with quarks.

1. He also computed their charge-to-mass ratio.

2. Isospin used to be called isotopic spin.

The existence of hadrons made up of three identical quarks with the same spin³, combined with the impossibility for fermions to have the exact same quantum numbers, indicated that quarks had additional degrees of freedom. This additional degree of freedom was called “color” and identified as the charge of the strong nuclear force, which binds quarks together and confines them into hadrons. The full description of the strong interaction is then provided by a theory called Quantum Chromodynamics (QCD) [5].

At the same time as progress was made to explain the hadron properties and their strong interactions, the weak nuclear force and the electromagnetic interaction were better understood. Electromagnetism alone was already successfully explained by the Quantum Electrodynamics (QED) theory developed by Feynman [6] and others. Low energy phenomena caused by the weak nuclear force were described by Fermi theory, but the theory broke down at high energy because the interaction rate grows ever higher with no upper bound. The introduction of new massive mediator bosons would solve this divergent interaction rate problem, but there seemed to be no consistent way to assign a mass to the bosons. This problem was solved with the Brout-Englert-Higgs mechanism [7, 8], where spontaneous symmetry breaking assigns a mass to the otherwise massless mediator bosons. At high energy the electromagnetic and weak forces are replaced by the electroweak forces [9]: the photon and the Z and W bosons are mixtures of the electroweak bosons.

The combination of QCD and electroweak theory is what is now called the Standard Model of particle physics.

1.2 A more detailed view of the Standard Model

1.2.1 Standard Model particles and forces

The SM provides a description of the electromagnetic force, and also of the strong and weak nuclear forces. No attempt is made to describe gravity. The model is based on a gauge symmetry, which means that its Lagrangian is invariant not only under global symmetry transformations, but also under local transformations, which may depend on spacetime coordinates. In order to enforce invariance under local transformations, it is necessary to introduce vector fields with special transformation properties into the Lagrangian. These vector fields have associated particles called gauge bosons and mediate the forces described in the SM.

The model is also explicitly Lorentz-invariant, i.e. invariant under rotations and boosts from the symmetry group of special relativity.

The SM particle content consists of three generations of fermions, of the previously mentioned gauge bosons, and of the Brout-Englert-Higgs boson (“H boson” in the rest of this thesis), see table 1.1 for a summary. All particles have corresponding antiparticles, except for the photon and the Z and H bosons which are their own antiparticles. Particles and antiparticles have the same mass and lifetime but opposite charges. Particles are classified according to their properties, some of which are:

- Their **flavor**, which is a quantum number conserved by the strong and electromagnetic interactions, but is sometimes violated by the weak force. Conventionally, each quark in the second and third generation is assigned its own flavor (strangeness, charm, bottomness and topness). For the first generation, another quantum number is used, isospin. For leptons, each generation is assigned a separate flavor or lepton number, e.g. one would

3. For example, the Ω^- is composed of three strange valence quarks with parallel spins.

compute the electron lepton number of a group of particles by summing the number of electrons and electron neutrinos and subtracting the number of anti-electrons and electron anti-neutrinos. Separate numbers are calculated for the second and third lepton generations by replacing “electron” by “muon” and “tau” in the previous sentence. All three lepton numbers would be conserved in the SM if neutrinos were massless.

- Their **electric charge**, which determines how the particle is affected by the electromagnetic force.
- Their **color charge**, which determines the particle’s interactions via the strong force. Since differently colored particles cannot be observed individually and distinguished, we only specify whether a particle possesses a color charge or if it is color neutral.
- Their **spin**, which determines how the particle transforms under rotations. Spin can take any positive integer or half-integer value, corresponding to a specific representation of the rotation group, and also to the particle’s intrinsic angular momentum in reduced Planck constant (\hbar) units. The spin-statistics theorem relates a particle’s spin to the statistical equations it obeys: all particles with integer spin are bosons while those with half-integer spin are fermions. Bosons follow Bose-Einstein statistics which allow many identical particles to occupy the same quantum state, while fermions follow Fermi-Dirac statistics and are bound by the Pauli exclusion principle, which states that no two identical fermions can occupy the same quantum state. Fundamental particles are only known to exist in three spin varieties: 0 for the H boson, 1/2 for fermions, and 1 for gauge bosons. Those three spin varieties correspond to three different representations of the rotation group: the scalar representation (spin 0), the spinor representation (spin 1/2) and the vector representation (spin 1). This is why the H boson is called a scalar boson, fermions are described by spinors and gauge bosons are called vector bosons.
- Their **mass**, which determines the relation between the energy and the momentum of a freely propagating particle. Massless particles always propagate at the speed of light c . Note that this thesis uses the “natural units” of particle physics (cf. appendix A), and thus masses are given in electron-Volts or multiples thereof.
- Their **lifetime**, which corresponds to the average time needed for the particle to decay. Some particles are stable and never decay.

SM fermions are divided into those that are affected by the strong force, called quarks, and those that are not, called leptons. First generation fermions, except neutrinos, are the building blocks of atoms. The two extra generations are very similar to the first one, each first generation fermion has two higher generation equivalents sharing most of their properties. The only differences are that the second and third generation fermions have different flavors, higher masses, and are all unstable, with varying lifetimes⁴.

The gauge bosons can be divided according to which force they mediate. The photon is massless and mediates the electromagnetic force, gluons are also massless and mediate the strong force, while the W^\pm and Z bosons are massive and are responsible for the weak interaction. The electromagnetic force is the only long-distance-acting force of the three, because the weak force is exponentially suppressed by its massive gauge bosons, and the strong force is confined into color-neutral hadrons⁵ and exponentially suppressed outside of them.

The three forces derive from the SM gauge symmetry group $(SU(3)_C \times SU(2)_L \times U(1)_Y)$ and its spontaneous symmetry breaking. $SU(3)_C$, where the “C” stands for “color”, is the strong interaction symmetry group. This group only acts on particles with color charge, i.e. quarks

4. Neutrinos are too light for any decay to be observed. Instead, we observe them oscillating between different flavors.

5. This hasn’t been mathematically proven but is strongly suspected by theorists.

	Fermions			Gauge bosons	H boson
	Generation I	Generation II	Generation III		
Quarks	u up 2/3 $\simeq 2$ MeV	c charm 2/3 1.3 GeV	t top 2/3 173 GeV	W^\pm W^\pm bosons ± 1 80 GeV	H H boson 0 125 GeV
	d down -1/3 $\simeq 4$ MeV	s strange -1/3 $\simeq 100$ MeV	b bottom -1/3 4.7 GeV	Z Z boson 0 91 GeV	
Leptons	e electron -1 511 keV	μ muon -1 106 MeV	τ tau -1 1.777 GeV	γ photon 0 0	
	ν_e electron neutrino 0 < 1 eV	ν_μ muon neutrino 0 < 1 eV	ν_τ tau neutrino 0 < 1 eV	g gluons 0 0	

Table 1.1 – Standard Model particles. Within each box, the first line shows the particle's symbol, the second its name, the third its electric charge (in multiples of the elementary charge $e \simeq 1,6 \cdot 10^{-19}$ C), and the last line its mass. There are eight different gluons and each of the six quark flavors exists in three different colors.

and the $SU(3)_C$ gauge bosons themselves, the gluons. The $SU(2)_L \times U(1)_Y$ group corresponds to the electroweak interactions. It is spontaneously broken to the electromagnetism symmetry group $U(1)_{e.m.}$ via the Brout-Englert-Higgs mechanism, on which more detail is given in section 1.2.2. Before symmetry breaking, the $U(1)_Y$ gauge boson is massless and acts on each particle according to the value of one of their quantum numbers, called weak hypercharge. The $SU(2)_L$ "weak isospin" gauge bosons are also massless and act only on fermions with left-handed chirality, hence the "L" subscript. The chirality of a fermion is well-defined for two-component Weyl spinors, which describe all fermions before the spontaneous symmetry breaking, as opposed to four-component Dirac spinors, which describe all massive SM fermions after symmetry breaking. The chirality of a Weyl spinor can be either left- or right-handed, and describes under which representation of the Lorentz group it transforms. A Dirac spinor is a mixture of one left-handed and one right-handed Weyl spinors.

A particle's electric charge Q , weak isospin T_3 and hypercharge Y are related by the following formula:

$$Q = T_3 + \frac{1}{2}Y \quad (1.1)$$

After symmetry breaking, three of the four $SU(2)_L \times U(1)_Y$ gauge bosons acquire a mass and become the weak force W^\pm and Z bosons, while the last one remains massless and becomes the photon.

1.2.2 Spontaneous symmetry breaking via the Brout-Englert-Higgs mechanism

The weak force massive gauge bosons pose a problem because writing an explicit mass term for them in the SM Lagrangian would violate unitarity, making some probabilities infinite (an obvious inconsistency). The Brout-Englert-Higgs mechanism is a way to give a mass to these bosons while preserving unitarity. As a bonus, it also allows fermions to acquire mass. Massive fermions are a mixture of left-handed and right-handed chiral states, but only left-handed chiral states are affected by the $SU(2)_L$ group and thus the weak isospin force. Before symmetry breaking, left-handed and right-handed fermions are all massless and form two distinct groups, it is only after symmetry breaking that they mix and the corresponding mixture acquires a mass.

Before symmetry breaking, the electroweak sector of the SM Lagrangian is:

$$\mathcal{L}_{EW} = \sum_{\psi} \bar{\psi} \gamma^{\mu} i D_{\mu} \psi - \frac{1}{4} W_a^{\mu\nu} W_{\mu\nu}^a - \frac{1}{4} B^{\mu\nu} B_{\mu\nu} \quad (1.2)$$

where:

- ψ is a fermion field. Left-handed fermions are $SU(2)_L$ doublets and right-handed ones are $SU(2)_L$ singlets.
- $\mu, \nu = 0, 1, 2, 3$ are spacetime coordinate indices, with 0 denoting time and 1-3 space. The sum over repeated indices is implied⁶.
- γ^{μ} are the Dirac matrices.
- D_{μ} is the electroweak covariant derivative, with $D_{\mu} = \partial_{\mu} + ig' \frac{1}{2} Y B_{\mu} + ig \frac{1}{2} \sigma_L^a W_{\mu}^a$
 - σ_L^a are the Pauli matrices, the "L" subscript indicates they only act on left-handed fermions.
 - Y is the generator of the $U(1)_Y$ group, it acts on fields according to their weak hypercharge.
 - B_{μ} and W_{μ}^a are respectively the $U(1)_Y$ ("weak hypercharge") and $SU(2)_L$ ("weak isospin") gauge fields.
 - g' and g are respectively the $U(1)_Y$ and $SU(2)_L$ coupling constants.
- $-\frac{1}{4} W_a^{\mu\nu} W_{\mu\nu}^a$ and $-\frac{1}{4} B^{\mu\nu} B_{\mu\nu}$ are the kinetic terms for respectively the weak isospin and weak hypercharge fields. Note that these gauge-invariant kinetic terms can only be written because the fields are massless.

For the Brout-Englert-Higgs mechanism to work, a complex scalar $SU(2)_L$ doublet Φ , with weak hypercharge $Y_{\Phi} = 1$, is needed. One can write it as a function of four real fields ϕ_a :

$$\Phi = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix} \quad (1.3)$$

The introduction of a scalar doublet implies the existence of a scalar Lagrangian:

$$\mathcal{L}_H = (D_{\mu} \Phi)^{\dagger} D_{\mu} \Phi - V(\Phi^{\dagger} \Phi) \quad (1.4)$$

where the two terms are:

- A kinetic term for Φ with the covariant derivative D_{μ} : $(D_{\mu} \Phi)^{\dagger} D_{\mu} \Phi$.

6. Appendix A gives more information on special relativity and mathematical conventions used in this thesis.

- A potential $V(\Phi^\dagger\Phi)$, depending on the gauge invariant scalar expression $\Phi^\dagger\Phi$. Then, the most general gauge-invariant expression for the potential is $V(\Phi^\dagger\Phi) = \mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2$, because higher powers of $\Phi^\dagger\Phi$ would lead to a breakdown of the theory at high energy, similarly to the way it happens to the Fermi theory of weak interactions. Note that λ must be positive for the potential to have a minimum.

We are interested in what happens at low energies, close to the vacuum. In a vacuum, the potential $V(\Phi^\dagger\Phi)$ is minimized. A necessary condition for the minimum point of the potential is that its derivatives need to be zero:

$$\begin{aligned} V(\Phi^\dagger\Phi) &= \mu^2 (\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2) + \lambda(\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2)^2 \\ 0 &\equiv \frac{\partial V}{\partial \phi_a} = 2\mu^2\phi_a + 4\lambda\phi_a (\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2), \quad a = 1, 2, 3, 4 \end{aligned} \quad (1.5)$$

If $\mu^2 > 0$, then there is only one minimum: all $\phi_a = 0$. The spontaneous symmetry breaking happens when $\mu^2 < 0$. In that case, the point where all $\phi_a = 0$ is a local maximum, so the minimum or the minima must be found elsewhere, where at least one of the ϕ_a is non-zero. There is an infinity of such minima, but by convention, ϕ_3 is the field chosen to have a non-zero value v at the minimum, while the values of all other ϕ_a are set to zero. v can be related to μ^2 and λ by replacing $\phi_3 = v$ and $\phi_1 = \phi_2 = \phi_4 = 0$ in equation 1.5:

$$\begin{aligned} 0 &= 2\mu^2 v + 4\lambda v(v^2) \\ 0 &= \mu^2 + 2\lambda v^2 \\ v &= \sqrt{\frac{-\mu^2}{2\lambda}} \end{aligned} \quad (1.6)$$

The value of Φ at its minimum is called its vacuum expectation value:

$$\langle \Phi \rangle = \begin{pmatrix} 0 \\ v \end{pmatrix} \quad (1.7)$$

From equations 1.4 and 1.6, the scalar sector of the SM Lagrangian can be rewritten like this⁷:

$$\mathcal{L}_H = \left| \left(i\partial_\mu - g'\frac{1}{2}YB_\mu - g\frac{1}{2}\sigma_L^a W_\mu^a \right) \Phi \right|^2 - \lambda (\Phi^\dagger\Phi - v^2)^2 \quad (1.8)$$

If the scalar doublet is rewritten as an expansion around its vacuum expectation value:

$$\Phi' \equiv \Phi - \langle \Phi \rangle = \Phi - \begin{pmatrix} 0 \\ v \end{pmatrix} \quad (1.9)$$

and we inject it in equation 1.8, then terms mixing v , three of the scalar ϕ_a fields, and the gauge fields B_μ and W_μ^a appear. With a suitable choice of gauge, these can be interpreted as mass terms, and the three ϕ_a fields are absorbed into a new definition of the gauge bosons. The fourth scalar degree of freedom corresponds to the H boson. There is a final complication in that the B_μ field and one of the three W_μ^a fields mix and produce the massive Z boson and the massless photon. The massive W^\pm bosons arise from the two other W_μ^a fields.

To summarize: at high enough energies, the difference between the fake minimum (all $\phi_a = 0$) and the true minimum ($\langle \Phi \rangle$) is irrelevant and interactions are mediated by massless

7. Ignoring the constant term λv^4 that does not change the underlying physics.

electroweak bosons. At lower energies, the difference becomes important and the electroweak interactions become the electromagnetic force mediated by the massless photon and the weak force mediated by the massive W^\pm and Z bosons. This is called spontaneous symmetry breaking.

Fermion masses come from their Yukawa interactions with the scalar doublet, which mixes left- and right-handed chiral states. The Yukawa Lagrangian is:

$$\mathcal{L}_{\text{Yukawa}} = -\bar{Q}_L^i \Gamma_u^{ij} \tilde{\Phi} u_R^j - \bar{Q}_L^i \Gamma_d^{ij} \Phi d_R^j - \bar{L}_L^i \Gamma_e^{ij} \Phi e_R^j + \text{h.c.} \quad (1.10)$$

where:

- the Q_L^i are the left-handed quark $\text{SU}(2)_L$ doublets (one for each of the three generations). A doublet is composed of one up-type quark and one down-type quark.
- the L_L^i are the left-handed lepton doublets, composed of one charged lepton and one neutrino.
- u_R^j , d_R^j and e_R^j are $\text{SU}(2)_L$ singlets, respectively the right-handed up-type quarks, down-type quarks and charged leptons. Note that there are no right-handed neutrinos.
- $\tilde{\Phi} = i\sigma_L^2 \Phi^*$, where Φ^* is the complex conjugate of Φ and $\sigma_L^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$ is the second Pauli matrix. The introduction of $\tilde{\Phi}$ is necessary because Φ would only allow interactions between the down-type quarks.
- Γ_u^{ij} , Γ_d^{ij} and Γ_e^{ij} are 3x3 complex matrices determining the fermion interactions with the scalar field.
- h.c. means hermitian conjugate.

If we do the substitution from 1.9 again, mass terms proportional to v and linking left- and right-handed fermions appear, as illustrated in equation 1.11 below. This is similar to what happens with the gauge bosons.

$$\begin{aligned} Q_L^i &= \begin{pmatrix} u_L^i \\ d_L^i \end{pmatrix}, \quad L_L^i = \begin{pmatrix} \nu_L^i \\ e_L^i \end{pmatrix} \\ \mathcal{L}_{\text{Yukawa}} &\ni -(\bar{u}_L^i \quad \bar{d}_L^i) \Gamma_u^{ij} \begin{pmatrix} v \\ 0 \end{pmatrix} u_R^j - (\bar{u}_L^i \quad \bar{d}_L^i) \Gamma_d^{ij} \begin{pmatrix} 0 \\ v \end{pmatrix} d_R^j - (\bar{\nu}_L^i \quad \bar{e}_L^i) \Gamma_e^{ij} \begin{pmatrix} 0 \\ v \end{pmatrix} e_R^j + \text{h.c.} \\ \mathcal{L}_{\text{Yukawa}} &\ni -v \cdot \bar{u}_L^i \Gamma_u^{ij} u_R^j - v \cdot \bar{d}_L^i \Gamma_d^{ij} d_R^j - v \cdot \bar{e}_L^i \Gamma_e^{ij} e_R^j + \text{h.c.} \end{aligned} \quad (1.11)$$

In this construction neutrinos remain massless. It is known by now that neutrinos do have masses, which is the only way to explain the flavor oscillation observed by e.g. Super-Kamiokande [10]. There is no widely-agreed way to include neutrino masses in the SM, there are many possibilities, but note that there is the straightforward solution of introducing right-handed neutrinos and adding a term similar to the up-type quark term, making use of $\tilde{\Phi}$.

1.2.3 The full Lagrangian of the Standard Model

The full particle content and interactions of the SM can be specified by writing its Lagrangian. Parts of it were already shown in section 1.2.2 in order to explain the Brout-Englert-Higgs mechanism. The Lagrangian is constructed by:

1. Introducing the fermion fields ($\psi = Q_L, L_L, u_R, d_R$ and e_R) and their kinetic terms.
2. Requiring Lorentz invariance.

3. Requiring the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge invariance of the fermion kinetic terms. They cannot be of the form $\bar{\psi}\gamma^\mu\partial_\mu\psi$, because that would not be invariant under the transformation $\psi \rightarrow U(x)\psi$ where $U(x)$ depends on spacetime coordinates. This necessitates the introduction of vector fields with special transformation properties, the gauge fields, and the partial derivative ∂_μ needs to be substituted by the covariant derivative D_μ which contains gauge field terms in a way that ensures gauge invariance. Kinetic terms for the gauge field also need to be added.
4. Introducing the $SU(2)_L$ scalar doublet needed for the Brout-Englert-Higgs mechanism, its potential (with a non-zero vacuum expectation value) and its kinetic term.
5. Including all possible renormalizable interaction terms. Any interaction term that can be written between any of the SM fields should be included, unless they are non-renormalizable. This last requirement is equivalent to disallowing any interaction terms where the coupling constant would have negative mass dimensions.

The full SM Lagrangian is then

$$\mathcal{L}_{\text{SM}} = \sum_{\psi} \bar{\psi}\gamma^\mu i D_\mu \psi + |i D_\mu \Phi|^2 - \lambda \left(\Phi^\dagger \Phi - v^2 \right)^2 + \mathcal{L}_{\text{Gauge}} + \mathcal{L}_{\text{Yukawa}} \quad (1.12)$$

where:

- D_μ is the covariant derivative of the full Lagrangian, with an extra term compared to the electroweak-only Lagrangian. $D_\mu = \partial_\mu + ig_s T^a G_\mu^a + ig \frac{1}{2} \sigma_L^a W_\mu^a + ig' \frac{1}{2} Y B_\mu$, where g_s is the strong force coupling constant, T^a are the generators of the $SU(3)_C$ group and G_μ^a the corresponding gauge fields.
- The kinetic terms for the gauge fields are in $\mathcal{L}_{\text{Gauge}} = -\frac{1}{4} G_a^{\mu\nu} G_{\mu\nu}^a - \frac{1}{4} W_a^{\mu\nu} W_{\mu\nu}^a - \frac{1}{4} B^{\mu\nu} B_{\mu\nu}$
- $\mathcal{L}_{\text{Yukawa}}$ is the same as in equation 1.10.

1.2.4 Why lepton flavor violating processes are forbidden in the Standard Model

The SM has a certain number of accidentally conserved quantities, which are not explicitly imposed as symmetries of the theory but which are conserved anyway as a consequence of other aspects of the model. Those accidental conservation laws are usually related to flavor quantum numbers. For example, one can assign a baryonic number $B = +1$ for baryons ($B = \frac{1}{3}$ for quarks) and $B = -1$ for antibaryons ($B = -\frac{1}{3}$ for antiquarks), and the total baryonic number of a collection of particles will not change under SM processes. The same happens if one introduces a leptonic number L , with $L = +1$ for leptons and $L = -1$ for antileptons. In fact, in the absence of neutrino masses, lepton numbers are conserved individually for each of the three generations, i.e. lepton flavor violating (LFV) processes are not allowed in the SM.

A comparison with the quark flavor sector is useful to understand why this accidental lepton flavor conservation occurs. There are three terms in the SM Lagrangian where quark flavors can mix: the weak force term in the covariant derivative and the two quark-scalar Yukawa interaction terms (one allowing up-type quarks to have masses, the other one for down-type quarks). The matrices Γ_u^{ij} and Γ_d^{ij} from equation 1.10 are generic complex matrices which can be diagonalized by a suitable biunitary transformation, i.e. by multiplying them on the left and the right by two different unitary matrices. Four independent (unitary) changes of basis can be performed for the left and right-handed down-type and up-type quarks, as illustrated in equation 1.13 below:

$$\begin{aligned}
U_{u_R}^\dagger U_{u_R} &= U_{u_L}^\dagger U_{u_L} = U_{d_R}^\dagger U_{d_R} = U_{d_L}^\dagger U_{d_L} = 1 \\
u_R &\rightarrow u'_R = U_{u_R} u_R, \quad u_L \rightarrow u'_L = U_{u_L} u_L \\
d_R &\rightarrow d'_R = U_{d_R} d_R, \quad d_L \rightarrow d'_L = U_{d_L} d_L
\end{aligned} \tag{1.13}$$

The four unitary matrices U_{u_R} , U_{u_L} , U_{d_R} and U_{d_L} are independent and can be chosen in such a way that they diagonalize the Yukawa matrices Γ_u and Γ_d :

$$\begin{aligned}
\bar{u}_L \Gamma_u u_R &= \bar{u}_L U_{u_L}^\dagger U_{u_L} \Gamma_u U_{u_R}^\dagger U_{u_R} u_R = \bar{u}'_L D_u u'_R \\
\bar{d}_L \Gamma_d d_R &= \bar{d}_L U_{d_L}^\dagger U_{d_L} \Gamma_d U_{d_R}^\dagger U_{d_R} d_R = \bar{d}'_L D_d d'_R
\end{aligned} \tag{1.14}$$

where D_u and D_d are diagonal matrices. This diagonal form corresponds to the usual understanding of mass, non-diagonal mass terms would imply that particles mix into each other while they are freely propagating through space. However, this comes at the cost of making the charged weak interactions non-diagonal. First let us consider the part of the SM Lagrangian containing the weak interactions:

$$\mathcal{L}_{\text{SM}} \ni \bar{Q}_L \gamma^\mu i D_\mu Q_L \ni -\frac{g}{2} \bar{Q}_L \gamma^\mu \sigma_L^a W_\mu^a Q_L = -\frac{g}{2} (\bar{u}_L \quad \bar{d}_L) \gamma^\mu \sigma_L^a W_\mu^a \begin{pmatrix} u_L \\ d_L \end{pmatrix} \tag{1.15}$$

The W_μ^a can be rearranged into a neutral boson W_μ^3 which does not change fermion flavor, and two charged bosons W_μ^+ and W_μ^- which e.g. allow up-type quarks to decay into down-type quarks and vice-versa. The part with the charged bosons is the one that interests us. We can switch to the diagonal mass basis:

$$\begin{aligned}
\mathcal{L}_{\text{SM}} &\ni -g \bar{u}_L \gamma^\mu W_\mu^+ d_L + \text{h.c.} \\
\mathcal{L}_{\text{SM}} &\ni -g \bar{u}_L U_{u_L}^\dagger U_{u_L} \gamma^\mu W_\mu^+ U_{d_L}^\dagger U_{d_L} d_L + \text{h.c.} = -g \bar{u}'_L \gamma^\mu W_\mu^+ U_{u_L} U_{d_L}^\dagger d'_L + \text{h.c.}
\end{aligned} \tag{1.16}$$

In the most general case, $U_{u_L} U_{d_L}^\dagger$ is not a diagonal matrix, and thus the exchange of W^\pm bosons allows mixing between different quark generations. The mixing matrix is known as the Cabibbo-Kobayashi-Maskawa (CKM) matrix, and contains four physical parameters. Three of them are mixing angles between quark generations and the other is a complex phase which violates the charge-parity (CP) symmetry⁸ [11].

Note that for the CKM matrix to exist, U_{u_L} and U_{d_L} must be distinct. In the lepton sector, if no masses are assigned to the neutrinos, only the charged lepton mass matrix needs to be diagonalized. The same change of basis can then be used for the left-handed charged leptons (equivalent to u_L in the quark analogy) and the neutrinos (equivalent to d_L), and as a consequence there is no equivalent for the CKM matrix in the lepton sector, and no LFV process. Adding neutrino masses does allow LFV processes to occur. Indeed, neutrino masses were found to be non-zero after the observation of an LFV process: neutrino flavor oscillation. Neutrino flavor eigenstates ($\nu_{L e}$, $\nu_{L \mu}$, $\nu_{L \tau}$) and mass eigenstates ($\nu_{L 1}$, $\nu_{L 2}$, $\nu_{L 3}$) are related by the so-called Pontecorvo-Maki-Nakagawa-Sakata (PMNS) unitary matrix:

$$\begin{pmatrix} \nu_{L e} \\ \nu_{L \mu} \\ \nu_{L \tau} \end{pmatrix} = V_{\text{PMNS}} \begin{pmatrix} \nu_{L 1} \\ \nu_{L 2} \\ \nu_{L 3} \end{pmatrix} \tag{1.17}$$

8. The CP symmetry is a combination of the charge (C) symmetry, which changes a particle into its antiparticle, and the parity (P) symmetry, which changes a left-handed chiral state into a right-handed chiral state and vice-versa.

The most straightforward SM-like neutrino mass term would be to introduce right-handed neutrinos and add a “Dirac” mass term similar to that of up-type quarks. The charged LFV process $\mu \rightarrow e\gamma$ is then possible (cf. figure 1.1), and its branching ratio can be shown to be [12] :

$$\mathcal{B}(\mu \rightarrow e\gamma) = \frac{3\alpha}{32\pi} \left| \sum_i (V_{\text{PMNS}})_{\mu i}^* (V_{\text{PMNS}})_{ei} \frac{m_{\nu_i}^2}{m_W^2} \right|^2 \quad (1.18)$$

where $\alpha \sim \frac{1}{137}$ is the fine structure constant from quantum electrodynamics, $m_W = 80 \text{ GeV}$ is the mass of the W boson, and $(V_{\text{PMNS}})_{\mu i}$ and $(V_{\text{PMNS}})_{ei}$ are PMNS matrix elements. The neutrino masses m_{ν_i} have not been measured yet, though stringent upper limits have been obtained. A direct search by the KATRIN collaboration [13] in Tritium β decay yielded an upper limit of 1.1 eV on the absolute neutrino mass scale, at the 90 % confidence level (CL). Tighter limits come from astrophysics and cosmology: a survey of the photometric redshift of over 700 thousand galaxies [14] found an upper limit of 0.28 eV on the sum of all neutrino masses at the 95 % CL, while the analysis of the cosmic microwave background by the Planck collaboration [15] found an upper limit of 0.23 eV on the summed neutrino mass at the 95 % CL.




Figure 1.1 – One possible Feynman diagram (see section 1.3 for more details on Feynman diagrams) contributing to an LFV $\mu \rightarrow e\gamma$ decay. This decay can occur because neutrinos have masses and their diagonal mass basis is different from their diagonal weak interaction basis.

From equation (1.18), even if we ignore the PMNS matrix elements⁹ and consider that the sum of squared neutrino masses is $\sum_i m_{\nu_i}^2 = 1 \text{ eV}^2$, the branching ratio is

$$\mathcal{B}(\mu \rightarrow e\gamma) \sim \frac{3\alpha}{32\pi} \left| \frac{1^2}{(80 \cdot 10^9)^2} \right|^2 \sim 5 \cdot 10^{-48} \quad (1.19)$$

i.e. an unobservably small number. Other charged LFV processes are similarly suppressed beyond any hope of experimental observation. This is why the observation of charged LFV would demand an entirely novel explanation, and would be an undeniable sign of new physics.

1.3 Standard Model predictions and experimental tests

All SM interactions can in principle be derived from the Lagrangian given in equation 1.12. For the electromagnetic and weak interactions this can be done by perturbation theory: the quantity to be calculated can be written as an infinite sum, each term being proportional to an ever higher power of the coupling constant. When the coupling constant is sufficiently small, higher-order terms contribute less than lower-order terms¹⁰, and truncating the sum makes sense as an approximation. When the coupling constant is too high, such as in low energy

9. These numbers are always smaller than 1, because the PMNS matrix is unitary.

10. At least up to a point. Technically the infinite sum is an asymptotic series, and does not always converge.

QCD, each higher order term contributes more than the last, and perturbation theory does not work. Other techniques, such as lattice QCD, are used. However, for high enough energies, the coupling constant of the strong force is low enough that perturbation theory is feasible.

Feynman diagrams are a convenient graphical way to write the terms of the infinite sums discussed in the previous paragraphs. They are made up of points (vertices) and lines, and the lines are labeled according to which particle/field they correspond to. Each line carries a spin and an energy-momentum Lorentz 4-vector, total energy-momentum is conserved at each vertex. Here, by convention, initial state particles are written on the left-hand side of the diagram, and final state particles on the right. Each initial and final state particle needs to obey the equation relating its energy E , momentum p and mass m :

$$E^2 - p^2 = m^2 \quad (1.20)$$

However no such requirement is imposed on internal lines. The Feynman rules relate each vertex and line to a specific mathematical expression, and by multiplying all of them together, one of the terms in the perturbative series is obtained. Figure 1.2 shows examples of Feynman diagrams, the one on the left is a so-called tree-level diagram because it contains no loop. Complicated, high order Feynman diagrams allow interactions that do not explicitly appear in the Lagrangian to occur. For example, the photon has no explicit self-interaction, but diphoton scattering can happen via a one-loop diagram such as in figure 1.2 (right). The central diagram in the figure is an example of a loop diagram adding a correction term to electron-positron scattering.




Figure 1.2 – Feynman diagrams of (left) tree-level electron-positron scattering, (center) a one-loop contribution to electron-positron scattering, and (right) photon-photon scattering.

A certain number of SM parameters need to be experimentally measured, but once this is done, predictions can be formulated for any particle physics process. The Feynman diagrams described above are the first step of a complex procedure to derive simulations of experimental processes, on which more details are given in chapter 6. At the end one obtains a simulation tailored to a specific experiment predicting the distribution of events which should be observed in data. The SM passed many such experimental tests, from the discovery of gluons at DESY [16] to the observation of the H boson by the CMS [17] and ATLAS [18] collaborations at CERN, almost 50 years after its existence was first proposed. But there are still some experimental and theoretical questions which the SM is unable to answer.

1.4 Standard Model incompleteness and hints of new physics

One obvious way in which the SM is incomplete is that it does not include a description of gravity. Combining in a consistent way the current best description of particle physics (the SM) with the current best description of gravity (General Relativity, GR) remains an open problem. Physical phenomena are well-described by the SM in the high energy density weak curvature (i.e. weak gravity) regime, and by GR in the low energy density strong curvature regime, but it is unclear what happens when high energy density and strong curvature are combined. While it is possible to quantize gravity in a similar way as with the other fundamental interactions, the procedure is non-renormalizable [19]. Research is still ongoing to find a consistent theory of quantum gravity, and combining it with the SM.

One other deficiency of the SM is that it does not include neutrino masses, despite experiments showing that they do have non-zero masses. As mentioned before adding Yukawa interactions between the neutrinos and the H field would give neutrinos masses in the same way as the other fermions. This would imply the existence of right-handed neutrino fields, as of yet unobserved. There is also a possibility that neutrinos could be their own antiparticles and would then have a so-called Majorana mass.

Another puzzle is the so-called strong CP problem. As mentioned when discussing the quark mixing matrix, in section 1.2.4, the weak interactions violate the CP symmetry. Experimentally, this is the only source of CP violation that has been observed so far. The strong force was at first thought to conserve CP because unlike the weak force, it couples equally to left- and right-handed chiral states. However, a careful analysis of the topological properties of the SU(2) gauge symmetry group reveals that it is always possible to write a renormalizable CP-violating term in the Lagrangian [20]. As a consequence, any gauge theory including SU(2) as a subset, like the strong force with its SU(3) symmetry, must also allow CP-violating interactions. But no CP violation phenomena relating to the strong force have been observed, and the θ_{CP} phase associated to them is constrained to be extremely small, in particular because of the upper bounds on the neutron electric dipole moment [21].

CP violation implies that matter and antimatter do not behave in exactly the same way. If the early universe is assumed to contain matter and antimatter in equal amounts, then CP violation is a necessary ingredient for the matter-antimatter asymmetry observed in the present-day universe [22]. However, the CP violation in the SM is too small to account for the asymmetry observed today [23], so theories with new sources of CP violation are proposed to explain it.

There are other astrophysical observations which fit poorly with our current understanding of physics, the main one being dark matter. Dark matter's existence was first proposed in the 1930s in order to explain the anomalous velocity distribution of galaxies in galaxy clusters. The velocities were incompatible with the cluster mass inferred from its visible components, the galaxies moved too fast to be gravitationally bound. There had to be an unobserved, "dark" matter contributing to the mass of the cluster and explaining the velocity distribution [24]. Since then, more evidence has accumulated in favor of the dark matter hypothesis: the gravitational lensing caused by galaxies is stronger than it should be, and the structure of the temperature fluctuations in the cosmic microwave background (CMB) can only be explained by the presence of large amounts of matter which does not interact with photons or baryons [25]. Dark matter thus remains the leading candidate to explain these anomalies, but it could be anything from fuzzy dark matter, a field with a mass of just 10^{-20} eV [26], to primordial black holes weighing many solar masses [27]. If dark matter turns out to be a new kind of fundamental particle anywhere in that mass range, the SM will have to be modified in order to describe it.

Dark matter is one component of the Λ CDM model, which describes the universe at cosmological scale. It combines cold dark matter (CDM) and a non-zero positive cosmological constant (Λ). Λ is interpreted as the vacuum energy density of the universe, the energy density of empty space, when no particles are present. It acts as a repulsive force, pushing things away from each other, as opposed to the usual gravitational interaction between two matter fields, which is attractive. A non-zero positive Λ is necessary to explain the acceleration of the expansion of the universe, first observed in 1998 with a measurement of the distance-redshift relation in supernovae [28]. The SM does predict the existence of a vacuum energy, but any attempt to calculate it yields a value which is many orders of magnitude larger. The effective value of Λ can be set to the observed quantity by renormalization, but the large difference between the bare and effective values of Λ is sometimes considered a hint for new physics [29].

The bare and effective values of the H boson mass are also very different. This is often called the H boson mass hierarchy problem and comes from the fact that the quantum loop corrections to the mass of any fundamental scalar field depend quadratically on the energy scale used as a cutoff when regularizing the Lagrangian. The effective mass of the H boson, the one that can be measured by experiment, is the difference between its bare mass (a parameter in the Lagrangian) and the quantum corrections. With an energy cutoff at the Planck scale ($\sim 10^{19}$ GeV), the quantum corrections to the mass are enormous and need to be almost exactly canceled by the bare mass parameter in order to produce an effective mass many orders of magnitude smaller [30]. The hierarchy problem is thus ultimately connected to the vast difference between the Planck scale (the scale at which quantum gravity effects become important) and the electroweak scale (the scale of the masses of the W, Z and H bosons). There is the possibility that a careful study of the renormalization group equations will reveal that no new physics above the electroweak scale is required [31], but many new physics theories, some of which are explained in more detail in the next chapter, solve this hierarchy problem by removing the quadratic divergence of the mass of the H boson.

Chapter summary

During the 20th century, particle physics experiments demonstrated the existence of many new phenomena. Those were described by new theoretical models, which eventually coalesced into a consistent whole called the Standard Model of particle physics. Matter is described by three generations of fermions, forces are carried by vector bosons whose existence is derived from gauge invariance. Fermions and vector bosons acquire their masses via the Brout-Englert-Higgs mechanism. All possible interactions are included in the SM, as long as they are renormalizable, and respect Lorentz and gauge symmetry. Sometimes those three requirements, combined with the SM matter content, conspire to disallow certain kinds of interactions, producing accidental conservation laws. Lepton flavor conservation is an example of such a law. The observation of neutrino masses has already challenged this part of the SM, and there are many other open questions which motivate the search for new physics. Charged lepton flavor violation is a generic consequence of many new physics models.

Beyond the Standard Model, towards charged lepton flavor violation

There are many models of particle physics beyond the standard model, which address the field's outstanding questions in different ways. Many experimental searches for such new physics have been performed, whether in high energy colliders, lower energy particle physics experiments, or by the observation of cosmic rays. Lepton flavor non-conservation has already been observed in the neutrino sector, and this at the very least implies charged lepton flavor violation via quantum loops. Moreover, since lepton flavor is only accidentally conserved in the SM, modifications of the theory which do not explicitly impose the conservation of lepton flavor often predict tree-level charged LFV. Such signatures are ideally suited for a collider search at the energy frontier.

This chapter starts with discussions of new physics models containing charged LFV processes, with R-parity violating supersymmetry in section 2.1, quantum black holes in section 2.2, LFV heavy H bosons in section 2.3, and an overview of Z' models in section 2.4. The chapter closes with a description of the experimental state of the art in section 2.5.

2.1 R-parity violating supersymmetry

2.1.1 A short overview of supersymmetry theory

Supersymmetry (SUSY) is a class of extensions to the SM which postulates a new symmetry linking bosons and fermions to each other. In a world where SUSY is an unbroken symmetry, there would no longer be separate fermion and boson fields, but supermultiplets containing both bosons and fermions. Bosons and fermions belonging to the same supermultiplet are called superpartners. In practice, this is achieved by introducing new spinor operators Q and Q^\dagger turning bosons into their fermion superpartners and vice-versa. The anticommutation of Q with its hermitian conjugated operator yields a translation P^μ . Schematically [32], the following relations can be written:

$$\begin{aligned}
 Q |\text{boson}\rangle &= |\text{fermion}\rangle \\
 Q |\text{fermion}\rangle &= |\text{boson}\rangle \\
 \{Q, Q^\dagger\} &= P^\mu \\
 \{Q, Q\} &= \{Q^\dagger, Q^\dagger\} = 0 \\
 [Q, P^\mu] &= [Q^\dagger, P^\mu] = 0
 \end{aligned} \tag{2.1}$$

SUSY is thus an extension of the Poincaré spacetime symmetry group, which ordinarily contains only rotations, translations and boosts. The final equation implies that particles and their superpartners would have the same masses: from the commutation of Q and Q^\dagger with the momentum operators P^μ ¹, it follows that they also commute with the square momentum operator P^2 , whose eigenvalues are the particles' masses squared. Each SM particle would have its own superpartner: SM chiral fermions would be paired with complex scalars (called sfermions) and SM gauge bosons with chiral fermions (called gauginos)². But light superpartner particles have never been observed, so SUSY can only be compatible with experimental results if it is broken at some scale.

One of the main arguments for SUSY is that it provides a solution to the scalar mass hierarchy problem mentioned in chapter 1. The hierarchy problem comes from the fact that the loop corrections to the H boson mass are quadratic. With SUSY, the fermions have scalar superpartners which give equal but opposite sign loop contributions to the H boson mass. The same happens with SM bosons and their fermionic superpartners. In broken SUSY, the fermion and boson contributions no longer have equal magnitudes, and some fine-tuning of the H mass is necessary. But the positive and negative contributions are still similar enough as long as the splitting between superpartner masses is not too far above the H mass. If SUSY is to be a solution to the hierarchy problem, TeV-scale superpartners need to be found.

2.1.2 R-parity conservation and violation

Many fermionic numbers, such as baryon number B and lepton number L , are conserved in the SM. That is a problem for SUSY, because of its introduction of a large number of new scalar bosons, the superpartners of the SM fermions. The new scalar bosons could mediate interactions between SM particles which violate B or L at such high rates that they should already have been observed by experiment. This problem can be solved by introducing a new discrete Z_2 symmetry, R-parity. Ordinary SM particles are assigned $R_p = +1$ and their superpartners $R_p = -1$, and conservation of this quantity is all that is needed to disallow unwanted exchanges of scalar bosons and thus preserve B and L [33]. Another advantage of R-parity is that it provides for a natural dark matter candidate: the lightest supersymmetric particle (i.e. the lightest particle such as $R_p = -1$), which is stable since it cannot decay into anything without violating R-parity [34]. R-parity can be expressed as:

$$R_p = (-1)^{2S}(-1)^{3(B-L)} \quad (2.2)$$

where S is the particle's spin. Let us check whether this equation gives the desired R-parity values. R-parity is factorized into a spin component $(-1)^{2S}$ and a "matter parity" component $(-1)^{3(B-L)}$ which equals $+1$ for SM bosons and -1 for SM fermions. Bosons have integer spin ($(-1)^{2S} = +1$) while fermions have half-integer spin ($(-1)^{2S} = -1$), so for SM particles the spin and the matter parity factors are always equal, which guarantees $R_p = +1$. If sfermions are assigned the same B and L of their fermion superpartner, then their matter parity factor will be negative while their spin component is positive since they are bosons. On the other hand, the fermionic superpartners of the gauge and H bosons carry no B or L , hence a positive matter parity factor, while their half-integer spin produces a negative spin factor. As expected, R-parity is negative for all of the as-of-yet unobserved superpartners of the existing SM particles.

1. The momentum operators are closely related to translations, in a way described by Noether's theorem.

2. Note that this pairing happens before the electroweak spontaneous symmetry breaking explained in chapter 1. Chiral fermions are massless, and so are the gauge bosons here. This also means that the superpartner of the right-handed electron is different from the one of the left-handed electron.

The Minimal Supersymmetric Standard Model (MSSM) is usually constructed with R-parity conservation and its building blocks are similar to those of the ordinary SM:

1. Chiral fermion superfields.
2. Lorentz invariance.
3. $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge invariance, with gauge superfields.
4. Spontaneous electroweak symmetry breaking with two H doublet superfields. Two H doublets are necessary instead of one because of the need to avoid quantum anomalies³ and the coupling of fields with different chiralities in the Lagrangian [35].
5. Including all renormalizable and R-parity preserving interactions.
6. Some SUSY-breaking term(s).

The so-called superpotential contains bilinear and trilinear terms with H, quark and lepton superfields. From this superpotential the part of the MSSM Lagrangian with the H mass term and the trilinear interactions of the H superfield with quarks and leptons can be derived. The superpotential is written below [33]:

$$W_{MSSM} = \mu H_u H_d + \lambda_{ij}^e H_d L_i E_j^c + \lambda_{ij}^d H_d Q_i D_j^c - \lambda_{ij}^u H_u Q_i U_j^c \quad (2.3)$$

where:

- H_u and H_d are the two H $SU(2)_Y$ doublet superfields, including their higgsino superpartners.
- L_i and Q_i are respectively the lepton and the quark $SU(2)_Y$ doublet superfields, including their slepton and squark superpartners.
- E_j^c , D_j^c and U_j^c are $SU(2)_Y$ singlet superfields⁴, respectively for charged leptons, down-type quarks and up-type quarks.
- μ is a mass term for the H superfields, and λ_{ij}^e , λ_{ij}^d and λ_{ij}^u are coupling constants for the trilinear interactions.
- i and j are indices running over the three fermion generations.

The lepton superfields L_i and the H superfield H_d have the same gauge quantum numbers: they are both left-handed, have no color charge, and -1 weak hypercharge. So gauge symmetry allows H_d to be substituted by any L_i in equation 2.3, the only thing stopping such terms to be written is R-parity conservation. If R-parity violation (RPV) is allowed, then both bilinear and trilinear lepton number violating terms appear in the superpotential. Here we focus only on the trilinear terms :

$$W_{RPV} = \frac{1}{2} \lambda_{ijk} L_i L_j E_k^c + \lambda'_{ijk} L_i Q_j D_k^c \quad (2.4)$$

The λ'_{ijk} terms allow lepton and down-type quark superfields to mix and the λ_{ijk} terms allow LFV couplings. In a model where the tau sneutrino $\tilde{\nu}_\tau$ is the lightest supersymmetric particle (LSP), it could be created by quark-antiquark annihilation and then decay into differently-flavored charged leptons, a signal signature which can be observed in collider experiments. The same couplings contribute to loop diagrams for freely propagating neutrinos, and hence to neutrino masses, as shown in figure 2.1.

3. Anomalies are discussed in more detail in section 2.4.

4. The ^c superscript denotes conjugation, the convention is to always use left-handed superfields.




Figure 2.1 – Left: tau sneutrino production by $d\bar{d}$ annihilation via a non-zero λ'_{311} coupling, then decay into the $\mu\tau$ LFV final state through the λ_{323} coupling. Right: the same trilepton coupling λ_{323} contributes to neutrino masses.

2.2 Quantum black holes

2.2.1 Extra dimensions and the weakness of gravity

In particle physics, gravity is an exceptionally weak force when compared to the three SM forces. Its effects only become manifest at the Planck scale, $\sim 10^{19}$ GeV, they can otherwise safely be neglected at collider experiments where the highest energy scale probed is a few TeV. However, theories with extra compact dimensions can bring down the Planck energy scale, possibly even down to experimentally accessible energies. This happens because in these theories, gravity is not inherently weak when compared to the other forces, it only appears weak because while the three SM forces stay confined to the usual four-dimensional spacetime, gravity “leaks” into the extra dimensions. This is a proposed solution to the hierarchy problem, which is ultimately connected to the vast gulf between the electroweak symmetry breaking scale and the Planck mass.

By bringing quantum gravity effects to the TeV scale, these models raise the possibility of lepton flavor violation. As shown in the previous chapter, lepton flavor is accidentally conserved, it is not associated to any fundamental symmetry of the SM. Quantum gravity may well induce LFV phenomena.

The simplest model with extra dimensions is one with n extra flat compact dimensions of radius $\sim R$ [36]. For distances r much smaller than R , the gravitational potential $V(r)$ between two masses m_1 and m_2 is given by Gauss’s law:

$$V(r) \sim \frac{m_1 m_2}{M_{\text{Planck}}^{n+2}} \frac{1}{r^{n+1}}, \quad r \ll R \quad (2.5)$$

where M_{Planck} is the Planck mass. For distances much larger than R , gravity can no longer propagate equally in all directions, it can at most go to a distance of R in the extra dimensions. The potential then follows the modified law:

$$V(r) \sim \frac{m_1 m_2}{M_{\text{Planck}}^{n+2}} \frac{1}{R^n r}, \quad r \gg R \quad (2.6)$$

This expression is identical to the gravitational potential from which Newton’s law is derived. An effective four-dimensional Planck scale can be identified, $(M_{\text{Planck}}^{\text{eff}})^2 = M_{\text{Planck}}^{n+2} R^n$, it has the usual value of $\sim 10^{19}$ GeV, while the real M_{Planck} could be at a much lower scale according to the size and the number of extra dimensions. This class of theories is called ADD (the initials of the authors of the seminal paper [36]) or Large Extra Dimensions (LED) because the extra dimensions would need to be much bigger than the Planck length scale in order to

bring quantum gravity effects to the TeV scale.

An alternative theory with a single small extra dimension has also been proposed [37], which is known as the Randall-Sundrum (RS) model. Instead of relying on the large size of the extra dimensions to bridge the gap between the electroweak and the Planck scales, the RS model introduces an extra dimension with an exponential “warp” factor. The usual flat four-dimensional spacetime metric ds^2 is given by:

$$ds^2 = \eta_{\mu\nu} dx^\mu dx^\nu \quad (2.7)$$

where $\eta_{\mu\nu}$ is a diagonal matrix of signature $(-1, +1, +1, +1)$ or $(+1, -1, -1, -1)$ depending on which convention is used. In the RS model, an extra dimension of size r_c is added to the metric, and the four-dimensional part is multiplied by an exponential function of the coordinate along the extra dimension:

$$ds^2 = e^{-2kr_c\phi} \eta_{\mu\nu} dx^\mu dx^\nu + r_c^2 d\phi^2 \quad (2.8)$$

where k is an energy scale close to the Planck mass, and $\phi \in [0, \pi]$ is the coordinate along the extra dimension. This five-dimensional spacetime would be bounded by two “branes”, a visible brane at $\phi = \pi$ corresponds to the usual four-dimensional spacetime, and a hidden one at $\phi = 0$. Fields with Planck scale masses in the hidden brane have effective masses at the TeV scale on the visible brane, thanks to the exponential warp factor. Gravity appears weak in the visible brane because the warped metric concentrates most of the graviton field near the hidden brane.

2.2.2 From extra dimensions to quantum black holes

Black holes are expected to be produced in two-particle collisions where the center-of-mass energy is above the Planck scale. These energies would normally be many orders of magnitude above what current experiments can access. However, we have just shown that in order to explain the weakness of gravity, models with large extra dimensions (ADD), or with a warped spacetime metric (RS), bring the fundamental Planck scale to only a few TeV. The observation of the decay products of small black holes in collider experiments is thus one of the main consequences of theories with extra dimensions. The production cross section σ of these black holes can be approximated from geometric considerations in the classical regime:

$$\sigma \sim \pi r_h^2(E = \sqrt{s}) \quad (2.9)$$

where $r_h(\sqrt{s})$ is the black hole horizon radius for a collision with center-of-mass energy \sqrt{s} . More careful estimates yield a similar expression for the classical cross section [38], and, assuming the impact parameter between the two colliding particles is large enough, quantum gravity corrections to this cross section should be small [39]. The geometric cross section from equation 2.9 can then be taken as a reasonable approximation of the black hole production cross section.

Black holes with masses much higher than the fundamental Planck mass can decay into multiple particles, through Hawking radiation [40]. Although they obey all local conservation laws, they otherwise decay with equal probability into all particles, in high-multiplicity final states [41]. However, black hole production is an inelastic process, where a lot of the collision energy may be radiated away by gravitational waves instead of being captured in the black hole. This reduces by many orders of magnitude the black hole production rates in collider experiments [42]. It is thus likely that if any black hole is observed at such an experiment, its mass will be close to the production threshold, and it will not be a semi-classical thermal

object. The alternative is the production of quantum black holes (QBHs), which decay mostly to two-particle final states [43]. Unlike semi-classical black holes, their decay properties would depend on the details of quantum gravity, which can be a problem as predictions are harder, but also an opportunity, because their observation would give more insight into the theory of quantum gravity. One of the most dramatic signatures of QBH would be decays that violate global quantum numbers, such as lepton flavor. Those would need to be explicitly forbidden by a new symmetry in order not to occur. LFV final states are thus an interesting way to probe quantum gravity at the TeV scale.

The specific details of the true theory of quantum gravity are certainly important, but one can nevertheless attempt a model-independent description of QBHs with a few general considerations. QBHs resemble more closely heavy resonances than semi-classical black holes. They can be characterized by their mass, spin, and gauge charges, including color⁵ and electric charge [44]. Conservation of gauge charges is assumed but not Lorentz invariance. It is possible that QBH couplings to highly offshell perturbative modes (e.g. exchange of virtual QBH between low energy particles) are suppressed, and therefore processes violating global quantum numbers can occur at the TeV scale without conflicting with low energy experimental constraints. Thanks to the universal nature of the gravitational coupling, it is expected that color-singlet and electrically neutral QBH decay equally into lepton pairs, whether flavor-diagonal or not. Their production cross section would depend on the threshold mass (linked to the Planck scale), the number of extra spatial dimensions n , and what specific extra dimensional model is used.

2.3 Lepton flavor violating heavy H bosons

As explained in detail in chapter 1, the scalar sector of the SM was introduced in order to provide a consistent gauge theory description of the electroweak force, and to allow fermions to acquire non-zero masses. A single complex scalar $SU(2)_L$ doublet is all that is needed to accomplish those two goals. But the scalar sector of new physics models could be more complicated, with e.g. an extra scalar doublet. The rich mass spectrum of models with two scalar doublets allows for new sources of CP violation, which is why they feature both in many proposed solutions to the strong CP problem [20] and in baryogenesis models explaining the matter-antimatter asymmetry of the universe [23]. And, as mentioned in section 2.1, a second H doublet is necessary in SUSY (see also the detailed overview in [45]). A generic feature of models with two scalar doublets is that they allow LFV processes [35].

In two scalar doublets models, there are two complex scalar doublets with non-zero vacuum expectation values:

$$\langle \Phi_1 \rangle = \begin{pmatrix} 0 \\ v_1 \end{pmatrix}, \langle \Phi_2 \rangle = \begin{pmatrix} 0 \\ v_2 \end{pmatrix} \quad (2.10)$$

where Φ_1 and Φ_2 are the two scalar doublets and v_1 and v_2 the two vacuum expectation values. There are eight degrees of freedom associated with these two doublets, and after electroweak symmetry breaking three of them get absorbed into the definition of the now massive W^\pm and Z bosons, as in the SM, and the other degrees of freedom correspond to physical scalar particles. Unlike the SM, there are five remaining degrees of freedom instead of just one. They correspond to one charged scalar H^\pm , two neutral scalars h (the lighter one) and H (the heavier one), and one pseudoscalar A .

5. This is not in contradiction with QCD confinement because the size of the QBH would be much smaller than that of a typical hadron.

LFV originates from the fact that leptons can now have Yukawa couplings with two scalar doublets instead of one⁶:

$$\mathcal{L}_{\text{Yukawa}} \ni -\bar{L}_L^i \Gamma_{e1}^{ij} \Phi_1 e_R^j - \bar{L}_L^i \Gamma_{e2}^{ij} \Phi_2 e_R^j \quad (2.11)$$

and as a consequence the charged lepton mass matrix M_e^{ij} is a linear combination of the two Yukawa coupling matrices:

$$M_e^{ij} = \Gamma_{e1}^{ij} v_1 + \Gamma_{e2}^{ij} v_2 \quad (2.12)$$

Diagonalizing the mass matrix no longer guarantees the diagonalization of the Yukawa couplings, i.e. the H boson(s) can couple differently-flavored leptons to each other. This discussion also holds for the quark sector, where flavor changing neutral currents (FCNC) are predicted. However, the model has to be constructed carefully in order not to clash with experimental constraints: if the flavor-changing Yukawa couplings are large, then the mass of the scalar mediator has to be large as well. For example, if the heaviest fermions dictate the scale of the flavor-violating Yukawa couplings, then the mass of the heavy neutral scalar mediating these interactions has to be higher than 150 TeV [35], much higher than what can be probed in direct searches in current collider experiments. Lighter scalars are possible if the couplings are smaller, such as in models with flavor symmetries [46, 47]. Those would allow the existence of a heavy H boson with LFV couplings and a mass under 1 TeV.

2.4 Lepton flavor violating heavy Z' bosons

A generic feature of many theories extending the SM is the presence of an additional $U(1)'$ gauge symmetry, with its associated Z' gauge boson. This can happen in theories with extra dimensions, Grand Unified Theories (GUT), “little Higgs” models, and string theory.

Theories with extra dimensions have already been discussed in section 2.2 in the context of quantum black holes. Then, we assumed that the SM forces were confined to the usual four-dimensional spacetime, while gravity could propagate into the extra dimensions. However, if SM gauge bosons are also allowed to propagate into the extra dimensions, this could create high mass excitations of those gauge bosons [48]. The couplings of the excitations need not be the same as those of the original bosons, if quarks and leptons are confined to different areas of the higher-dimensional space [49].

Grand Unified Theories are based on the approximate convergence of the values of the three SM gauge couplings when they are extrapolated to an energy scale of $\sim 10^{16}$ GeV. They embed the SM gauge symmetry group $G_{\text{SM}} \equiv SU(3)_C \times SU(2)_L \times U(1)_Y$ into a larger group with a single gauge coupling, such as $SU(5)$ [50], $SO(10)$ [51] or even E_6 [52]. The larger symmetry is then spontaneously broken into the SM group, and the single gauge coupling splits into three independent couplings whose values diverge from each other at low energy scales. Lie groups, such as those making up the SM gauge symmetry, can be characterized by their rank, which is the maximum number of simultaneously diagonalizable generators of the group [53]. Spontaneous symmetry breaking can replace a symmetry group with one of lower or equal rank, but not one with a higher rank. The SM has rank 4, which means any GUT must have rank at least 4. A GUT based on a higher rank symmetry, such as $SO(10)$ (rank 5) or E_6 (rank 6), raises the possibility that the larger group could break into $G_{\text{SM}} \times U(1)'$ (rank 5). The earliest and simplest versions of GUT were disproven by the non-observation of proton decay [54], but it is possible to construct models which do not contradict experiments [55].

6. Object definitions are given in chapter 1.

“**Little Higgs**” models propose a solution to the little mass hierarchy problem. The solution to the H boson mass hierarchy problem was expected to be found at the electroweak scale, at ~ 200 GeV. New physics at the TeV scale would still pose a hierarchy problem, but a much smaller one. The purpose of little Higgs models is to solve this small hierarchy problem, bridging the gap between the TeV and electroweak scales. Like in SUSY, the one-loop quadratically divergent contributions to the H mass are cancelled by equal but opposite sign quadratically divergent contributions. Unlike SUSY, these cancellations happen between particles with the same statistics: SM boson contributions are cancelled by new bosons, the SM top contribution is cancelled by a new fermion [56]. In little Higgs models, the H fields are Goldstone bosons arising from a global symmetry breaking at an energy scale Λ_S higher than the electroweak scale. At the electroweak scale, they acquire a mass, which remains light thanks to the approximate global symmetry. In an explicit construction [57], the global symmetry is $SU(5)$, and it contains a locally gauged subgroup $[SU(2)_1 \times U(1)_1] \times [SU(2)_2 \times U(1)_2]$. At the energy scale Λ_S , the global symmetry is broken ($SU(5) \rightarrow SO(5)$) and at the same time the gauge subgroup is broken into the electroweak group $SU(2)_L \times U(1)_Y$. At this stage the SM W and B bosons remain massless (they acquire a mass at the electroweak scale), but the other gauge bosons acquire a mass close to Λ_S . And so there is a new massive gauge boson arising from a $U(1)'$ symmetry.

String theory is an attempt to unify gravity and the SM forces into a single theoretical framework. It is often combined with SUSY, in which case it is called “superstring theory”. Realistic superstring-inspired models must contain the SM gauge group, the three fermion families and two H doublets, i.e. the particle content of the minimal supersymmetric standard model. They also usually contain extra $U(1)'$ symmetries, whose corresponding Z' bosons could have masses at the TeV scale without too much fine-tuning [58].

Constraints from anomaly cancelations

One must be careful that the addition of an extra $U(1)'$ symmetry to the SM does not add any gauge anomalies, otherwise the theory would be inconsistent. Quantum anomalies are one-loop diagrams giving non-renormalizable infinite contributions to cross section calculations. They arise when a symmetry which holds in the classical version of a theory no longer holds at the quantum level [59]. The presence of chiral fermions in the SM causes gauge anomalies, which must cancel out in order to make the theory consistent. As a consequence, the following equations must be true [60]:

$$\begin{aligned} \sum_{f \in \text{fermions}} Y_f &= 0, & \sum_{f \in \text{fermions}} Y_f^3 &= 0, \\ \sum_{f \in (\text{anti})\text{quarks}} Y_f &= 0, & \sum_{f \in SU(2)_L \text{ doublets}} Y_f &= 0 \end{aligned} \quad (2.13)$$

where Y_f is the weak hypercharge of the fermion f . Factors of three have to be included for the quark colors and factors of two for the $SU(2)_L$ doublets. When substituting the SM charges in the left-hand side of all four equations, all contributions do cancel out to zero. An additional $U(1)'$ symmetry would impose further anomaly-cancelling equations. If the $U(1)'$ charge is called Q_2 , then all equations in 2.13 must hold when replacing Y by Q_2 . In addition, there will be equations mixing Y and Q_2 :

$$\sum_{f \in \text{fermions}} Y_f Q_{2f}^2 = 0, \quad \sum_{f \in \text{fermions}} Y_f^2 Q_{2f} = 0 \quad (2.14)$$

These equations cannot be satisfied by the SM fermions alone if the $U(1)'$ charges are the same for all three generations. This would imply the existence of additional “exotic” fermions, or preferential coupling to some generations [60]. As a consequence, FCNC and LFV phenomena are generally present in many of these models.

2.5 The experimental state of the art in charged lepton flavor violation searches

2.5.1 Low energy searches for lepton flavor violation

There has long been an interest in experimental searches for charged lepton flavor violation. As discussed in chapter 1, lepton flavor conservation is an accidental symmetry of the SM with zero neutrino masses. Even a straightforward SM extension with non-zero neutrino masses, which allows neutrino flavor oscillation, predicts unobservably small charged LFV. However, the status of lepton flavor as an accidental symmetry means that models extending the SM generally allow charged LFV processes, unless they specifically impose a symmetry guaranteeing at least approximate lepton flavor conservation. The combination of extremely small SM cross sections and the potential for much higher cross sections in new physics scenarios explains why many experimental searches for charged LFV have been performed over the last decades, as can be seen in figure 2.2.




Figure 2.2 – Best experimental upper limits (at the 90% confidence level) on various charged LFV processes over the years. The expected results of upcoming experiments are included. Collider searches at the high energy frontier are not shown here. For muon or tau decay searches, the upper limit is on the branching ratio of the looked-for decay mode. For $\mu \rightarrow e$ conversion, the upper limit is on the ratio of converted muons to captured muons. Image copied from [61]. The image dates from 2018, all limits shown later are projections.

Figure 2.2 also shows that muon-electron mixing is the most strongly constrained charged LFV process. The possible signatures include:

- $\mu^- N \rightarrow e^- N$, where a muon turns into an electron after interacting with a heavy nucleus⁷ N . Muons stopped by a heavy nucleus quickly form a muonic atom with the nucleus. The muon can then either decay in orbit ($\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$) or be captured by the nucleus, muon

7. The presence of the nucleus is necessary in order to conserve energy and momentum.

to electron conversion is a hypothetical third option. The current best experimental upper limit on the ratio of $\mu^- \rightarrow e^-$ conversion to muon capture comes from the SINDRUM collaboration [62], which set it at 7×10^{-13} at the 90% confidence level (CL).

- $\mu^+ \rightarrow e^+\gamma$, an (anti)muon decaying to an (anti)electron and a photon. This search is usually done with antimuons rather than muons because the latter can be captured by nuclei instead of decaying. The usual way in which a muon decays to an electron includes neutrinos in the final state in order to balance lepton flavor; if a photon replaces the neutrinos in the final state the process is LFV. The strongest upper limit on the branching ratio⁸ of $\mu^+ \rightarrow e^+\gamma$ is 4.2×10^{-13} (90% CL), and comes from the full dataset of the MEG experiment [63].
- $\mu^+ \rightarrow e^+e^+e^-$, an antimuon decaying to two positrons and one electron. The current best upper limit on the branching ratio of such a process comes from a relatively old paper from the SINDRUM collaboration [64], which set it at 1.0×10^{-12} (90% CL).

LFV processes involving taus are much less constrained, because it is harder to produce a large number of taus than to create a high-intensity muon beam. Nevertheless, there are experimental constraints on the following LFV tau decays:

- $\tau \rightarrow l\gamma$, where a tau decays to a photon and a light lepton $l = e, \mu$. The Belle experiment upper limits at the 90% CL on the branching ratios are 4.5×10^{-8} for $\tau^- \rightarrow \mu^-\gamma$ and 1.2×10^{-7} for $\tau^- \rightarrow e^-\gamma$ [65]. The BABAR collaboration also set limits at 90% CL: $\mathcal{B}(\tau \rightarrow \mu\gamma) < 4.4 \times 10^{-8}$ and $\mathcal{B}(\tau \rightarrow e\gamma) < 3.3 \times 10^{-8}$ [66].
- $\tau^- \rightarrow \mu^- \mu^- \mu^+$, a tau decaying into three muons. Upper limits on the branching ratio at 90% CL have been set by Belle, BABAR and LHCb, they are respectively 2.1×10^{-8} [67], 3.3×10^{-8} [68] and 4.6×10^{-8} [69]. ATLAS and CMS have also searched for this decay, though their 90% CL upper limits on the branching ratio are less strong: 8.0×10^{-8} for CMS using 2016 data [70] and 3.8×10^{-7} for ATLAS using 2012 data [71].

2.5.2 Lepton universality

Apart from neutrino oscillations, there is another related, but not identical, experimental challenge to the lepton flavor sector of the SM: significant deviations from lepton universality have been observed in decays of B hadrons (i.e. hadrons with a b valence quark⁹) in BABAR, Belle and LHCb. Lepton universality is the fact that the three lepton generations have identical gauge couplings in the SM. The three charged leptons have the same couplings to the W and Z bosons and the photon, the same is true for the three neutrinos. The only difference between the three generations comes from their Yukawa couplings to the scalar doublet, which also gives them different masses. Violations of lepton universality do not necessarily imply charged LFV, but they are often present in the same new physics models.

The experiments probing lepton universality measure ratios of the event rates of similar processes which differ only in the flavor of the leptons. The advantage of this method is that systematic errors which occur both in the numerator and the denominator of the ratio cancel out. Two sets of ratios show tensions between the SM and experimental measurements:

- $R_{D^{(*)}} = \frac{\mathcal{B}(B \rightarrow D^{(*)} \tau \nu_\tau)}{\mathcal{B}(B \rightarrow D^{(*)} l \nu_l)}$ where B is any of B^+ (valence quarks: $u\bar{b}$), B^- ($\bar{u}b$), B^0 ($d\bar{b}$) or \bar{B}^0 ($\bar{d}b$); D is any of D^+ ($c\bar{d}$), D^- ($\bar{c}d$), D^0 ($c\bar{u}$) or \bar{D}^0 ($\bar{c}u$); l can be either an electron or a muon. The D mesons can also come in their excited versions D^* .

8. The branching ratio of a particle X decaying into the final state Y is the number of $X \rightarrow Y$ decays divided by the total number of X decays to any final state.

9. The valence quarks determine the overall flavor of a hadron. Hadrons may contain other quark-antiquark pairs, called sea quarks, which have no net flavor contribution.

- $R_{K^{(*)}} = \frac{\mathcal{B}(B \rightarrow K^{(*)} \mu^+ \mu^-)}{\mathcal{B}(B \rightarrow K^{(*)} e^+ e^-)}$ where K can be K^+ ($u\bar{s}$), K^- ($\bar{u}s$), K^0 ($d\bar{s}$) or \bar{K}^0 ($\bar{d}s$). The ratio can also be computed with the excited K^* mesons. These B decays are FCNC and are only possible at the one-loop level in the SM, they are thus less frequent than the decays in the previous bullet point.

Combining LHCb, BABAR and Belle measured values of R_D and R_{D^*} ¹⁰, there is a 3.6-3.8 σ tension with the SM, depending on the specific choice of the SM R_{D^*} prediction [72]. The predictions and the experimental results in the R_D - R_{D^*} plane are shown in figure 2.3. The BABAR and Belle measurements of R_K and R_{K^*} are compatible with the SM due to large experimental relative uncertainties, as shown in figure 2.4. LHCb does not compute $R_{K^{(*)}}$ directly, instead it calculates a double ratio to the resonant decay $\mathcal{B}(B \rightarrow K^{(*)} J/\psi \rightarrow K^{(*)} l^+ l^-)$. Electrons in LHCb have high bremsstrahlung rates due to their high energies, and even with dedicated procedures the efficiency of the reconstruction of electrons is significantly lower than the muon reconstruction efficiency. The double ratio method allows most of these differences to cancel out, and enables LHCb to keep uncertainties low. The LHCb results show a 2.1-2.6 σ tension with the SM [72], as illustrated in figure 2.4.

These discrepancies found in experimental tests of lepton universality are another reason to investigate flavor physics in more detail.




Figure 2.3 – SM predictions and latest experimental values from BABAR, Belle and LHCb for the R_D and R_{D^*} ratios, showing a tension of more than 3 σ [73].

10. In LHCb these ratios are computed using only muons.




Figure 2.4 – SM predictions and recent experimental values from BABAR, Belle and LHCb for the R_K and R_{K^*} ratios. The BABAR and Belle results are in agreement with the SM, while the LHCb values are more than 2σ away [72].

2.5.3 High energy searches for lepton flavor violation

The non-observation of charged LFV in muon/tau decay and muon conversion, discussed in section 2.5.1, constrains the LFV couplings of heavier particles too. If a heavy neutral boson X could decay into both e^+e^- and μ^+e^- , then it could also mediate the decays $\mu^+ \rightarrow e^+e^+e^-$ and $\mu^+ \rightarrow e^+\gamma$, as shown in figure 2.5.




Figure 2.5 – If a hypothetical boson X had LFV couplings in addition to flavor conserving couplings, then it could mediate LFV muon decays, such as $\mu^+ \rightarrow e^+e^+e^-$ (left) and $\mu^+ \rightarrow e^+\gamma$ (right).

In general, an effective non-renormalizable charged LFV (CLFV) Lagrangian can be written by integrating out heavy degrees of freedom. The effective Lagrangian is valid up to an energy scale Λ , at which point it needs to be replaced by a new Lagrangian specifying the heavy degrees of freedom. Low energy precision constraints on CLFV can be translated into constraints on the heavy scale Λ , which can be directly probed at high energy colliders if it is not too large. In general, the effective Lagrangian will contain many terms, but for the purposes of illustrating the links between different searches for CLFV physics, only two terms need to be considered. For example [74]:

$$\mathcal{L}_{\text{CLFV}} = \frac{m_\mu}{(\kappa + 1)\Lambda^2} \bar{\mu}_R \sigma_{\mu\nu} e_L F^{\mu\nu} + \frac{\kappa}{(\kappa + 1)\Lambda^2} \bar{\mu}_L \gamma_\mu e_L (\bar{u}_L \gamma^\mu u_L + \bar{d}_L \gamma^\mu d_L) + h.c. \quad (2.15)$$

where:

- μ , e , u and d are respectively the muon, electron, up quark and down quark fields.
- m_μ is the muon mass.
- $F^{\mu\nu}$ is the photon field strength, which features in the photon gauge kinetic term $-\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$.
- $\sigma^{\mu\nu} = \frac{i}{2}[\gamma^\mu, \gamma^\nu]$.
- κ is a dimensionless parameter determining the relative size of the two terms.

The first term of the effective Lagrangian dominates when $\kappa \ll 1$, it directly mediates $\mu \rightarrow e\gamma$ decay and mediates $\mu \rightarrow eee$ decay and $\mu \rightarrow e$ conversion at the next-to-leading order level. The second term, on the other hand, dominates when $\kappa \gg 1$, and it mediates $\mu \rightarrow e$ conversion at leading order and $\mu \rightarrow e\gamma$ and $\mu \rightarrow eee$ decays with next-to-leading order terms. The specifics of the Lagrangian can determine which of the charged LFV processes has a favored cross section, which is why it is useful to search separately for all of them. The CLFV experiments already mentioned in this section probe energy scales Λ up to ~ 1000 TeV in the case of e - μ mixing. New physics could of course exist below that scale, but its CLFV e - μ couplings would necessarily be tiny [74]. An exception to this rule could be quantum black holes, whose couplings to highly offshell perturbative modes are speculated to be heavily suppressed (cf. section 2.2).

The energy scale Λ at which the heavy CLFV degrees of freedom become relevant is directly probed at high energy collider experiments. Searches for CLFV decays of both already known and hypothetical heavy states have been performed. In the case of known particles, there is more interest in decays involving tau leptons, given the extremely strong constraints on e - μ mixing. As of the writing of this thesis, these are the latest CLFV searches at high energy:

- CLFV decays of the H boson. The LHC experiments CMS and ATLAS¹¹ have the latest limits at the 95% CL: ATLAS found $\mathcal{B}(H \rightarrow e\tau) < 0.47\%$ and $\mathcal{B}(H \rightarrow \mu\tau) < 0.28\%$ [75] while CMS found $\mathcal{B}(H \rightarrow e\tau) < 0.61\%$ and $\mathcal{B}(H \rightarrow \mu\tau) < 0.25\%$ [76].
- CLFV decays of the Z boson. The latest ATLAS search set $\mathcal{B}(Z \rightarrow e\tau) < 5.8 \times 10^{-5}$ and $\mathcal{B}(Z \rightarrow \mu\tau) < 1.3 \times 10^{-5}$ at the 95% CL [77]. CMS has only searched in the $e\mu$ channel, finding $\mathcal{B}(Z \rightarrow e\mu) < 7.3 \times 10^{-7}$ at the 95% CL [78], while low energy experiments imply $\mathcal{B}(Z \rightarrow e\mu) < 5 \times 10^{-13}$ [79].
- Neutral heavy H boson with CLFV decays. The CMS search, in the $e\tau$ and $\mu\tau$ final states, did not find any deviation from the SM [80]. The upper limits at the 95% CL for the product of the production cross section and the LFV branching fraction of heavy H bosons with masses in the 200-900 GeV range vary from 94.1 to 2.3 fb and 51.9 to 1.6 fb for respectively the $e\tau$ and $\mu\tau$ final states.
- New resonance or quantum black hole with CLFV decays. ATLAS performed such a search with data collected in 2015-16 on all three possible final states, and found no deviation from the SM, with heavy LFV Z' bosons excluded at the 95% CL for masses below 4.5, 3.7, and 3.5 TeV for respectively the $e\mu$, $e\tau$ and $\mu\tau$ final states [81]. CMS searched only in the $e\mu$ final state, using data collected in 2016, and also found no deviation from SM expectations¹²: heavy LFV Z' bosons are excluded for masses up to 4.4 TeV [82].

The searches for CLFV in high energy experiments have so far yielded no significant deviation from SM expectations, much like the lower energy experiments. But there are still some gaps to be covered. The work presented in this thesis is the first CMS search for a heavy state with $e\tau$ or $\mu\tau$ decays, and the first search of its kind taking advantage of the full LHC 2016-18 data-taking campaign¹³, in which protons collided at a center-of-mass energy of 13 TeV, the highest ever reached in a particle collider.

Chapter summary

The status of lepton flavor conservation as an accidental symmetry of the SM means that theories extending the SM often predict LFV phenomena. This chapter presented many such theories which predict the existence of new heavy neutral states with charged LFV decays: R-parity violating SUSY and its tau sneutrino, theories with extra dimensions and their TeV-scale quantum black holes, two scalar doublet models and their heavy H bosons, and finally an overview of some SM extensions including new $U(1)'$ gauge symmetries and the associated Z' bosons. Experiments have probed these theoretical possibilities, looking for rare muon and tau decays, muon to electron conversions, CLFV decays of known SM neutral bosons, and also new heavy states with CLFV decays. No significant deviation from SM expectations has been found so far. The work presented in this thesis is a novel search for heavy neutral states with CLFV decays, complementing a previous CMS search with more final states analyzed, and more data.

11. More details on LHC and its experiments are given in chapter 3.

12. I was already part of the CMS analysis team searching for CLFV at high energy when this paper was published. However, I was working mostly on the tau channels.

13. More details about LHC data-taking in chapter 3.

Part II

Experimental hardware and software

The Large Hadron Collider and the Compact Muon Solenoid

The search for charged lepton flavor violation presented in this thesis would not be possible without a machine to accelerate and collide particles at high energy and a detector to collect the collision data. The collider in question is the Large Hadron Collider (LHC), described in section 3.1. The detector is called Compact Muon Solenoid (CMS), and it is described along its many subdetectors in section 3.2.

3.1 The Large Hadron Collider

The LHC is a particle collider operated by the European Organization for Nuclear Research (CERN). CERN's main site is located in Meyrin, Switzerland, near Geneva. Figure 3.1 shows the complex of particle accelerators run by CERN, which occupy a large area in the French-Swiss border region and provide particle beams for a wide range of experiments, such as nuclear and atomic physics experiments using the ISOLDE facility [83], analyses of the properties of antimatter using the Antiproton Decelerator [84], or studies of rare kaon decays with the NA62 experiment [85]. The LHC is the largest accelerator at CERN, and indeed, in the world.

3.1.1 Physical layout of the LHC

The LHC is an approximately circular machine installed in the 26.7 km tunnel which was previously occupied by the CERN Large Electron Positron collider (LEP). The LHC's approximate circle is actually made of eight arcs and eight straight sections, the latter of which can be used to install experiments or utilities. Two vacuum rings go around the machine, and host the two particle beams rotating in opposite directions. The particles can be either protons or heavy ions such as Pb^{82+} , and data taking campaigns have involved proton-proton, ion-ion and proton-ion collisions. The rest of this chapter focuses on proton-proton collisions, which are the only ones relevant for this thesis.

Radiofrequency (RF) cavities accelerate the beams, and in the arc sections dipole magnets keep the beams in the quasi-circular trajectory, while quadrupole magnets keep them focused in the transverse plane.

The eight straight sections are connected to the ground level by access shafts. Going clockwise around the machine, they host the following equipment [87]:

- Point 1 is the closest to the CERN Meyrin site. The ATLAS experiment [88] is located there. Together with CMS, ATLAS is one of the two big, general purpose detectors built to probe both proton-proton and ion-ion collisions.

CERN's Accelerator Complex

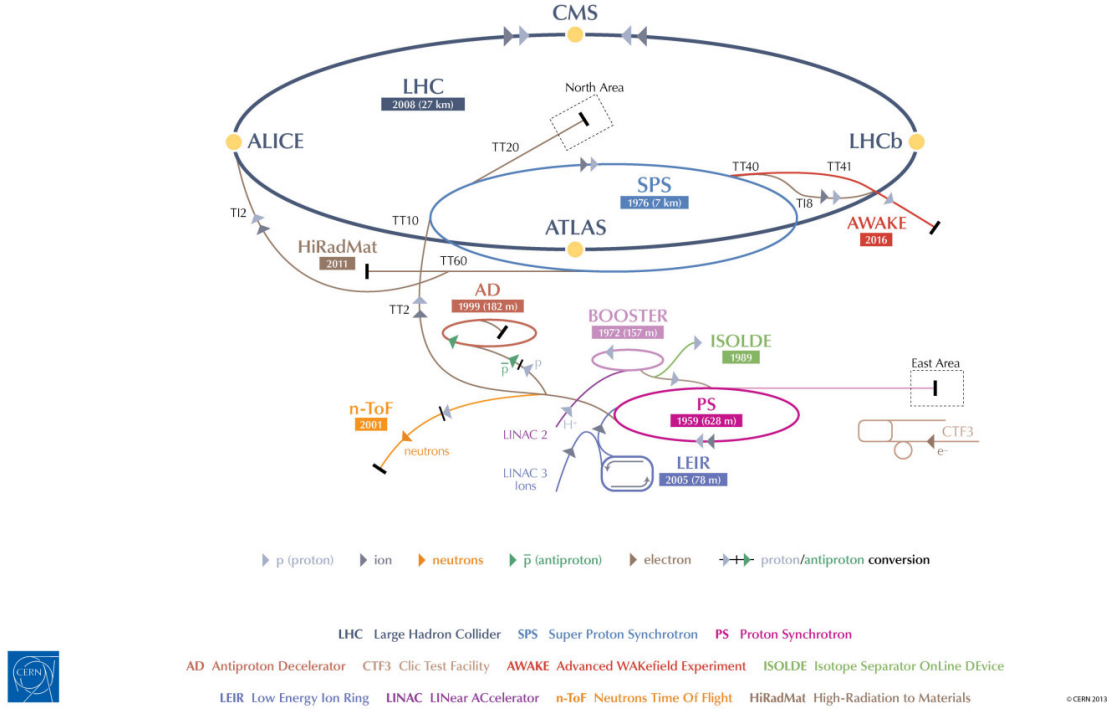


Figure 3.1 – Sketch of the CERN accelerator complex [86].

- Point 2 hosts the ALICE detector, which focuses on heavy ion collisions and the strong interaction sector of the SM at extreme values of energy density and temperature, where matter forms a quark-gluon plasma [89]. Point 2 is also the site of the injection line for the clockwise-rotating beam (as seen from above).
- Point 3 contains a collimation system which cleans the momentum spread of both beams. Particles with a high momentum offset compared to the rest of the beam are scattered and dumped into collimators.
- In Point 4 there are two superconducting RF systems, one for each beam [90]. An oscillating electric field is created in each cavity, the shape and the frequency of each cavity were chosen in such a way that a proton arriving at the right time and with the right energy is not accelerated. Protons with different energies, or arriving early or late, are decelerated or accelerated. This guarantees that protons stay into well-separated bunches, and that they stay close to the desired energy [91]. The desired energy can be increased over time, which is what allows the beams to go from the injection energy to the collision energy. There are also transverse damping and feedback systems at Point 4, which can correct injection errors and some instabilities in the transverse plane, and can also amplify transverse oscillations to facilitate beam measurements [92].
- Point 5 is the location of the Compact Muon Solenoid (CMS) detector [93], which collected the data analyzed in this thesis. CMS is described in more detail in section 3.2.
- Point 6 contains the beam dump systems, which function independently for each beam. This is where the beams are deflected and safely taken out of the LHC, in case of an emergency (e.g. there is a risk the beams could damage the machine) or a planned interruption.

- Point 7 is where a second collimation system is located. This one scatters particles with large transverse oscillating amplitudes.
- Point 8 houses the LHCb detector, which focuses on precision measurements of CP violation and rare decays of B hadrons [94]. The injection system for the counter-clockwise rotating beam is located nearby.

3.1.2 Particle beams in the LHC

The two most important parameters of a collider, which determine what physical phenomena it can probe, are its center-of-mass energy and its instantaneous luminosity. The LHC was designed to accelerate two counter-rotating proton beams to 7 TeV each, and then collide them at a center-of-mass energy of 14 TeV. Such high energies have not been reached yet, as of the writing of this thesis the highest center-of-mass energy reached was 13 TeV, or 6.5 TeV per beam.

A particle's energy E at rest is given by Einstein's famous formula:

$$E = mc^2 \quad (3.1)$$

where m is the particle's mass and c the speed of light in a vacuum. For a moving particle, the formula becomes:

$$E = \gamma mc^2 \quad (3.2)$$

γ is known as the Lorentz factor and depends on the particle speed (cf. appendix A). A proton's rest energy is about 938 MeV, a 6.5 TeV proton in the LHC has thus $\gamma \sim 6900$, which means its speed is very close to the speed of light.

The instantaneous luminosity determines the event rates of all physical processes at the LHC. Higher values of the instantaneous luminosity allow rarer processes, with smaller cross sections, to be observed. For a specific process with a cross section σ , the event rate $\frac{dN}{dt}$ is given by the following expression:

$$\frac{dN}{dt} = L\sigma \quad (3.3)$$

where L is the collider's instantaneous luminosity [87], which has units of inverse area per unit of time (e.g. $\text{cm}^{-2}\text{s}^{-1}$). L depends only on the beam parameters, it is higher when there are more protons in each beam, and lower when the proton bunches are wider at the collision point. Assuming the two beams are equal and that there are equal numbers of protons per bunch, the equation for the instantaneous luminosity is:

$$L = \frac{N_p^2 N_b f_{rev}}{4\pi\epsilon_n\beta^*} F \quad (3.4)$$

The meaning of the symbols on the right-hand side and their nominal LHC values are [87]:

- $N_b = 2808$ is the number of proton bunches in each beam.
- $N_p = 1.15 \times 10^{11}$ is the number of protons per bunch.
- $f_{rev} \sim 11.2$ kHz is the revolution frequency of a particle in the LHC, calculated by dividing the speed of light by the circumference of the machine.
- $\epsilon_n \sim 3.75$ μm is the normalized transverse beam emittance. It is proportional to the average beam width in the machine, and depends on the initial conditions of the beam production before it even gets to the LHC, and also on the beam energy E ($\epsilon_n \propto \frac{1}{\sqrt{E}}$).

- $\beta^* \sim 0.55$ m is the amplitude function at the collision point. The amplitude function β describes the variation of the actual beam width along the machine. Machine “optics” (using magnets) can trade a small β^* at the collision point, which is desirable to increase luminosity, for a larger β further away from the collision point. $\beta^* = 0.55$ m is the nominal LHC value for the high luminosity experiments CMS and ATLAS, but operators managed to squeeze it as low as 0.33 m in 2017-18 operations [95].
- F is a geometric factor translating the fact that the beams collide at an angle, and not head-on. In general $F \leq 1$, it would equal one if the collisions were exactly head-on.

The nominal peak instantaneous luminosity of the LHC is $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, but this value was increased to $1.9 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in 2017, mostly by reducing the transverse emittance and β^* [95].

However important its peak value may be, luminosity does not stay constant over a data-taking run. If no particular action is taken, luminosity decays exponentially as bunches collide and take protons out of the beams¹. Luminosity can also decrease due to an increase in emittance, which can be caused by particles scattering on residual gas (the beam pipe vacuum is not perfect), beam-beam interactions, intrabeam scattering and RF noise. This increase in emittance can be partially compensated by the radiation damping systems. Taking all of these effects into account, the luminosity lifetime² is about 15 h [87].

The hours-long lifetime of instantaneous luminosity means that LHC beams do not stay forever in the machine. In normal operating conditions, data-taking cycles follow each other as beams are repeatedly injected into the machine, then made to collide, and finally dumped when their luminosity is too low. Each cycle starts with the injection of test beams into the collider. These beams are low intensity, or have a low number of bunches, in order to stay below the machine damage threshold. They allow LHC operators to check if the machine is ready for higher intensity beams. Then the high intensity “physics” beams can be injected. This is a lengthy process, as the two entire 26.7 km LHC rings have to be filled, with up to 2808 bunches per ring. After that, the two beams are accelerated until they reach the collision energy, 6.5 TeV per beam during LHC Run 2. Then β^* is decreased so that the two beams are as narrow as possible at the collision point. If everything is under control, “stable beams” are declared and the collection of data suitable for physics analyses starts. When the luminosity gets too low the beam is dumped, and the machine is ramped down so that later a new beam can be injected. The smooth operation of the LHC, as just described, requires a sequence of complex tasks which have to be carried out in strict order. Each step of the cycle is categorized by a “beam mode” [96]. Beam modes for emergency stops are also included.

3.1.3 The LHC injection chain

Before a beam gets into the LHC, it needs to be accelerated to a minimum energy, and it needs the correct bunch structure. During nominal operations, proton beams injected into the LHC are supposed to have 2808 bunches separated by ~ 25 ns intervals, 1.15×10^{11} protons per bunch, and an energy of 450 GeV per proton. There is also a tight upper limit on the beams’ transverse emittance, which has to fit into the aperture of the LHC superconducting magnets. A long injection chain, including many different accelerators, handles the beams before sending

1. However, sometimes machine operators practice a so-called “luminosity levelling”, where they progressively squeeze the beams at the collision point to prevent natural luminosity decay. This is done when maximally squeezing the beams at the beginning of a run would produce a higher than desired luminosity. Instead, a lower maximum luminosity is achieved, but over a longer period of time. Nevertheless, this levelling usually cannot last for a whole data-taking run, and the luminosity eventually starts decreasing.

2. After which about a $\frac{1}{e}$ fraction of the initial luminosity remains.

them to the LHC.

Protons start their journey to the LHC in a hydrogen bottle, from which hydrogen atoms are extracted and ionized. The resulting protons are pre-accelerated to a kinetic energy of 750 keV and sent to a linear accelerator, Linac2. Linac2 accelerates protons by bunches up to a kinetic energy of 50 MeV, and then a transmission line sends them to the Proton Synchrotron Booster (PSB), a circular accelerator with four rings [97]. The PSB accelerates the proton bunches to 1.4 GeV and then injects them into the Proton Synchrotron (PS), a larger circular machine, which increases the protons' energy to 25 GeV. The PS also splits the bunches multiple times in order to produce an LHC bunch train with 25 ns spacing, while keeping the transverse emittance and the longitudinal size of the bunches under control [87]. The bunches are then sent to an even larger circular machine, the Super Proton Synchrotron, which accelerates them to 450 GeV. The bunches are then extracted from the SPS and injected into the LHC via two 2.8 km transfer lines [98]. Multiple SPS cycles are necessary in order to completely fill up the two LHC rings.

This injection chain was not built from scratch just for the LHC, it repurposed older accelerators built for other reasons. The SPS used to be the largest accelerator at CERN, it enabled the discovery of the W [99, 100] and Z [101, 102] bosons when operated as a proton-antiproton collider. The SPS was then used as a LEP injector before being adapted as an injector for the LHC. Similarly, at the time of its construction the PS was the largest accelerator at CERN, before it was repurposed as an injector to the SPS. Along with the Linac2 and the PSB, it had to undergo major upgrades to become LHC-ready [97].

The LHC is also not the final destination for all beams going through the accelerators mentioned in this section. For example, the SPS provides beams for the NA62 experiment, and ISOLDE uses proton bunches from the PSB.

3.1.4 Past and future operations timeline

The LHC has taken, since the year 2010, large quantities of data at increasingly higher center-of-mass energies and instantaneous luminosities. Physics data-taking was supposed to start earlier, but an accident occurred during the commissioning of the machine, in September 2008. An electrical arc developed between two magnets, triggering a chain reaction which led to the quench of several superconducting magnets and vacuum degradation in both of the beam pipes and in some insulating vacuum barriers. Some magnets were displaced and others even knocked off their supports [103]. The machine had to be shut down while the damage was repaired and operators investigated what happened and what could be done to prevent any further incidents. Commissioning with circulating beams started again in November 2009 [104].

When data-taking started in 2010, operators were cautious in order to protect the machine. The proton-proton collision center-of-mass energy (\sqrt{s}) was 7 TeV, and the instantaneous luminosity was progressively ramped up to a maximum value³ of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, by increasing the number of colliding bunches from 1 to 348 and by squeezing the beams at the interaction point [105]. The data delivered by the LHC in 2010 corresponds to an integrated luminosity of $44.2 \text{ pb}^{-1} = 44.2 \times 10^{36} \text{ cm}^{-2}$ (1 barn = 1 b is a unit of area equal to $1 \times 10^{-24} \text{ cm}^2$). 2010 marked the beginning of LHC Run 1, which lasted until the end of 2012. The integrated luminosity delivered in 2011 and 2012 was much greater than in 2010, respectively 6.1 fb^{-1} and 23.3 fb^{-1} , as shown in figure 3.2 (left). Data in 2011 was collected at $\sqrt{s} = 7 \text{ TeV}$, increasing to 8 TeV in 2012.

3. This value applies to the two high-luminosity experiments, CMS and ATLAS.

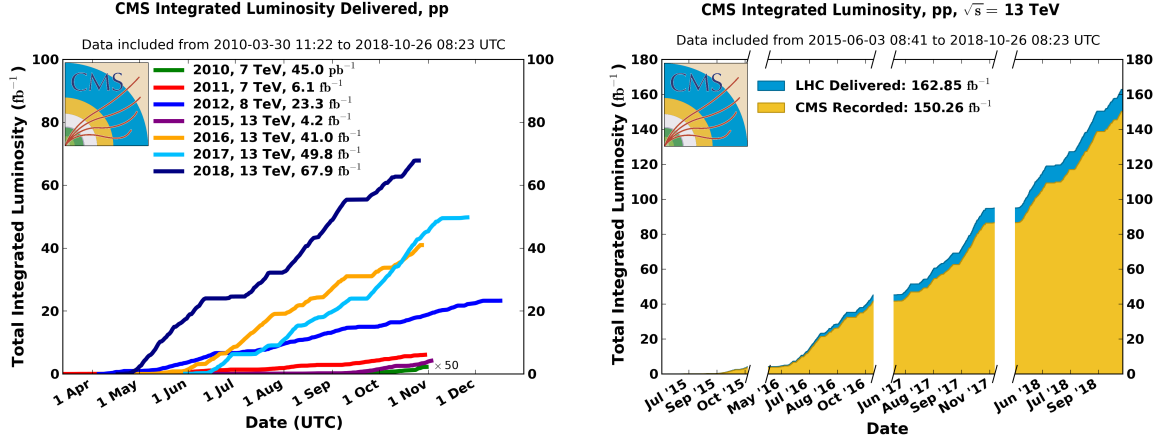


Figure 3.2 – (Left) Integrated luminosity delivered by the LHC, shown separately for each year of Run 1 and Run 2. (Right) Integrated luminosity over all of Run 2, in blue what was delivered by LHC at the CMS interaction point and in yellow what CMS recorded [106].

A planned shutdown in 2013-2015 (Long Shutdown 1 - LS1) was used for maintenance and preparing the machine and the detectors for operations at a higher center-of-mass energy, $\sqrt{s} = 13 \text{ TeV}$, during Run 2. Run 2 lasted from 2015 to 2018, with the LHC delivering increasingly higher luminosities year after year. 2015 was a commissioning year, and unexpected problems in the cryogenic system for the superconducting CMS magnet complicated operations. Many machine shutdowns were necessary in order to maintain the cooling system but even so CMS had to collect some data without any magnetic field [107]. As explained in section 3.2, the magnetic field bends charged particles in a way which depends on their momenta, and this is precisely how these momenta are measured. All of this meant that the LHC only delivered 4.2 fb^{-1} of integrated luminosity in 2015, including data when the CMS magnet was down. The luminosity delivered in the other three years of Run 2, 2016-18, was much higher: 41.0 fb^{-1} in 2016, 49.8 fb^{-1} in 2017 and 67.9 fb^{-1} in 2018, as shown in figure 3.2 (left). Full Run 2 CMS analyses usually do not include data taken in 2015, because its relatively small luminosity is not worth the extra work. The analysis presented in part III of this thesis is no exception, it was performed on data collected in 2016-18.

The luminosity actually used for analyses, in CMS or any of the other detectors, is always lower than what is delivered by the LHC. Many factors are at play:

- There is the detector deadtime, which can be due to CMS subsystems not being ready, or trigger rules (if an event triggers the data acquisition too soon after another, it is ignored), or calibration [108]. Deadtime was typically 3.5% during normal operations in Run 2, sometimes rising to 5%⁴.
- Occasionally, there are subsystem problems resulting in detector downtime, during which no data is collected at all.
- The data which does get collected by CMS needs to go through quality checks. Only certified data, collected when all subdetectors are working normally, is used in analyses.

The right side of figure 3.2 shows the difference between the integrated luminosity delivered by the LHC and the one actually collected by CMS over the full Run 2. The data collected by CMS is shown after deadtime and downtime, but before quality checks.

4. Source: my personal experience when on shift at the CMS control room, and checking deadtime rates on the CMS OMS tool.

Run 2 was followed by a second Long Shutdown (LS2), starting at the end of 2018 and scheduled to continue until 2022, which will mark the beginning of Run 3. Run 3 is planned to last until 2024, operating at $\sqrt{s} = 13$ or 14 TeV, and its total integrated luminosity will be roughly equal to that of Run 2. After that, LS3 will start in order to upgrade the machine and the detectors for even higher instantaneous luminosities: this is the HL-LHC (High Luminosity LHC) program. The target is a peak instantaneous luminosity of $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and an integrated luminosity of 250 fb^{-1} per year, with the eventual goal of reaching 3000 fb^{-1} by 2036 [109]. Figure 3.3 shows the long-term plans for the LHC.

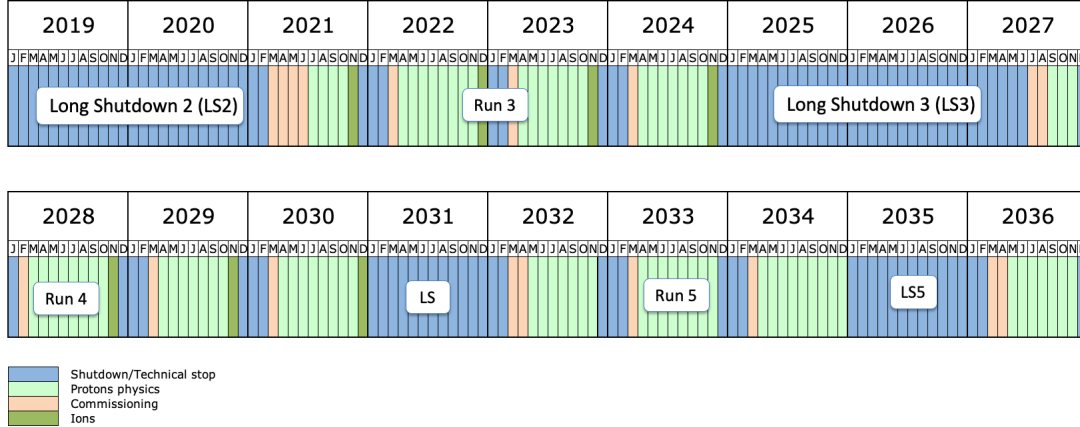


Figure 3.3 – LHC schedule until 2036 [110]. Note that due to the COVID-19 pandemic, Run 3 is now scheduled to start in early 2022, a few months delay compared to the late 2021 start shown in this figure. The HL-LHC operations start in Run 4.

3.2 The Compact Muon Solenoid

3.2.1 Overview of CMS and its coordinate system

The Compact Muon Solenoid (CMS) is a general purpose detector for proton-proton and heavy ion collisions located at LHC Point 5 in Cessy, France. As shown in figure 3.4, CMS is roughly cylindrical, and constructed in many subdetector layers around the beampipe, where the collisions happen. The full detector is 28.7 m long, has a diameter of 15.0 m and weighs about 14 000 t.

A silicon pixel and strip tracker is located closest to the interaction point, it tracks the paths of charged particles going through it. The tracker is surrounded by calorimeters, which measure particle energies by a destructive process. The innermost calorimeter is the electromagnetic one (ECAL), which handles photons and electrons⁵, and then comes the hadronic calorimeter (HCAL) for charged and neutral hadrons. A superconducting solenoid with an internal diameter of 6 m surrounds most of the HCAL and provides a 3.8 T magnetic field. The solenoid is kept cool at a temperature of $\sim 4.5 \text{ K}$ by a liquid helium refrigeration plant [93]. The outer part of the HCAL (HO) is located just outside of the solenoid in the central region of CMS. Gas ionization chambers are embedded in the magnet steel return yoke outside the solenoid and track the paths of charged particles able to punch through all of the detector, which are almost always muons. All subdetectors just described come in a central “barrel” section and

5. In this thesis, particles and antiparticles are referred to by the same name when the difference between them is not important. Unless noted otherwise “electrons” should be understood as “electrons and anti-electrons” (or “electrons and positrons”). See appendix A.




Figure 3.4 – Schematic view of the CMS detector [111], with all its subcomponents: the tracker, the calorimeters, the preshower, the solenoid, the steel return yoke and the gas ionization muon chambers.

two “endcaps”, one on either side of the barrel. A preshower detector sits in front of the ECAL endcaps, it was designed in order to help distinguish single photons from photon pairs coming from neutral pion decays. Beyond the endcaps, there are forward calorimeters providing an extended coverage.

All subdetectors are linked to the trigger system, which decides which events are kept in long-term storage and which are rejected and permanently lost. It plays a crucial role in reducing the data storage rate to a manageable level while keeping the most interesting events, as defined by the data analyzers. It is described in chapter 4.

The CMS coordinate system is centered on the nominal collision point, with an x-axis pointing to the center of the LHC, a roughly vertical y-axis pointing upwards, and a z-axis along the beampipe. The z-axis points in the anti-clockwise direction of the LHC as seen from above, a direction determined by the right-handedness of the coordinate system. The azimuthal angle ϕ is defined in the xy-plane with respect to the x-axis, and the polar angle θ is defined with respect to the z-axis. However, θ is not used to describe particle momenta, the pseudo-rapidity η is preferred:

$$\eta = -\ln \tan \left(\frac{\theta}{2} \right) \quad (3.5)$$

Pseudo-rapidity differences are conserved by Lorentz boosts in the z-axis, in the limit where the particle masses are negligible when compared to their momenta or energies⁶. There is no such conservation for the polar angle.

3.2.2 The silicon tracker

The CMS silicon tracker is the part of the detector located closest to the collision point, with a length of 5.8 m and a diameter of 2.5 m. Its purpose is to measure precisely and efficiently the trajectories of charged particles, as well as reconstructing secondary interaction

6. This is shown in appendix A.

vertices [93]. The tracker is located inside the radius of the solenoid which provides it with a homogeneous 3.8 T magnetic field. Charged particles are bent by this magnetic field according to their momenta, which can thus be measured by analyzing the curvature of the particle tracks.

Radiation resistance is a major constraint in the design of the tracker as its location very close to the interaction point implies a high density of particles flying through it. Another constraint is the so-called “material budget” of the detector: the more material there is in the way of a particle, the more undesirable interactions such as electron bremsstrahlung and photon conversion are likely. Some amount of material is of course unavoidable as it would otherwise be impossible to detect any particle, but the thickness of the silicon pixel cells and silicon strips, as well as the volume of readout cables, cooling systems and support structures, all have to be minimized.

Figure 3.5 shows a schematic cross section of the tracker as it was designed, including its many subcomponents. The innermost part of the tracker is a pixel detector consisting of many thin $100 \times 150 \mu\text{m}^2$ pixel cells, each with its own readout channel. This close to the interaction point, the use of pixel cells is necessary due to the extremely high rate of particle hits, and the excellent resolution the cells provide in all three dimensions also helps identify secondary vertices, e.g. those corresponding to tau decays. Until 2016, the pixel tracker was made of three cylindrical layers in the barrel and two disks in each of the endcaps. During the extended 2016-17 year-end shutdown, a detector upgrade [112] added a fourth layer to the barrel (moving the first layer closer to the beam line), as well as a third disk on each endcap. The upgrade included two extra features: an improvement of the readout chip, which would otherwise suffer from buffer overflow during operations at high instantaneous luminosity, and a reduction in the material budget by changing the cooling system, using a lighter mechanical support structure and relocating electronics boards to higher $|\eta|$.




Figure 3.5 – Schematic cross section of the CMS tracker design [93], with each line representing a detector module (double lines represent back-to-back modules with a stereo angle). The pixel detector was changed in the year-end LHC shutdown of 2016-17, including the addition of a fourth layer in the barrel and a third disk on both endcaps.

As we get further away from the interaction point, the area to be covered by the detector increases quadratically, and an ever-larger number of pixel cells would be needed. In order to keep the number of readout channels under control, the outermost tracker detector layers use

silicon strips, which are narrow in two dimensions but much longer in the third dimension. Some layers are equipped with two back-to-back strip detector modules with a stereo angle of 100 mrad in order to provide a measurement of the third coordinate. The strip tracker is divided into many sections:

- The Tracker Inner Barrel and Disks (TIB/TID), the innermost section, consists of four barrel layers and three disks at each end.
- The Tracker Outer Barrel (TOB), surrounds the TIB/TID in the radial direction and contains six detector layers.
- The Tracker Endcaps (TEC+ and TEC- according to their position along the z-axis) are composed of nine disks each and positioned on either side of the central part of the detector where the pixel, TIB/TID and TOB are located.

The strip detector provides a single point resolution (in the traverse direction) varying from 23 μm in the inner parts of the TIB, where strip density is the highest, to 53 μm in the outer rings of the TEC disks, where strip density is the lowest. The tracker as a whole covers the pseudo-rapidity range $|\eta| < 2.7$ after the 2016-17 pixel upgrade, but only up to 2.5 before that.

3.2.3 The electromagnetic calorimeter

The silicon tracker is completely surrounded by the electromagnetic calorimeter (ECAL), which is made of thousands of scintillating lead tungstate (PbWO_4) crystals, separated into a barrel and two endcap sections, a preshower detector being located in front of each of the endcaps. The crystals are connected to photodetectors which collect scintillating light: avalanche photodiodes (APDs) in the barrel and vacuum phototriodes (VPTs) in the endcaps [93].

The ECAL measures with high precision the energies of electrons and photons coming out of the tracker. This is achieved by placing a high density scintillating material, the aforementioned PbWO_4 , in the path of outgoing particles: any electron or photon entering the detector material triggers a series of cascading interactions (electron bremsstrahlung and photon pair production) and is effectively destroyed, producing an electromagnetic shower. For example, an electron entering the ECAL quickly emits a bremsstrahlung photon, which can in turn convert into an electron pair, while the original electron can emit a second bremsstrahlung photon, etc. The amount of matter traversed by an electron or photon between two interactions is characterized by the radiation length of the material, while the transverse size of the showers is characterized by the material's Molière radius. The PbWO_4 crystals are excited by the electromagnetic showers, the crystals then relax to a lower energy state by scintillating visible light, which is collected in the photodetectors. Note that particles other than electrons and photons are much less likely to trigger electromagnetic showers in the crystals chosen for the ECAL.

The ECAL is a homogeneous calorimeter, i.e. its entire volume is sensitive to incoming particles. PbWO_4 was chosen as the sensitive material thanks to its high density, short scintillating time⁷ and small radiation length and Molière radius [93]. However, the light output is relatively low and decreases with temperature [113]. Since the APD gain also drops with increasing temperature, a cooling system with precise temperature control is necessary.

The choice of photodetectors was not the same in the barrel and the endcaps due to their different radiation environments. The APDs, used in the barrel, have higher gain and quantum efficiency⁸ than the VPTs, but the latter are more radiation resistant and thus more suited for

7. 80% of the light is emitted in 25 ns, i.e. the time interval between two successive LHC bunches.

8. Higher quantum efficiency means an incoming scintillating photon is more likely to trigger a signal.

endcap operations [114].

The ECAL achieves a very high precision in energy measurements. Its energy resolution can be decomposed into a stochastic term S , a noise term N and a constant term C , with different energy dependences:

$$\left(\frac{\sigma_E}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2 \quad (3.6)$$

where E is the energy of the incoming particle and σ_E the error on the energy. Contributions to the stochastic term include fluctuations in the lateral shower spread, in the number of photoelectrons produced, in the gain of the APDs or VPTs, and in the energy deposited in the preshower detector [93]. The noise term is dominated by the electronics and the digitization process, with a small contribution from pileup, i.e. particles coming from secondary proton-proton interactions in the same event. The constant term dominates at high energy and is thus particularly relevant for the analysis presented in this thesis. The main issues contributing to it are the non-uniformity of the light collection along the longitudinal axes of the crystals and errors in channel-by-channel relative calibration, a small contribution comes from energy leakage from the back of the crystals. In test beams, without any magnetic fields or material in front of the ECAL, the energy resolution was measured to be [115]:

$$\left(\frac{\sigma_E}{E}\right)^2 = \left(\frac{2.8\%}{\sqrt{E(\text{GeV})}}\right)^2 + \left(\frac{12\%}{E(\text{GeV})}\right)^2 + (0.3\%)^2 \quad (3.7)$$

The energy resolution of electron/photon objects in real operating conditions is not as good, which is why it is also measured in the collisions data collected each year. In 2017, the energy resolution on electrons coming from Z boson decays (around 45 GeV of energy per electron) was about 1.5% in the most central part of the detector, increasing to $\sim 4.5\%$ in the endcaps [116].

A major challenge in keeping the ECAL well calibrated during data-taking is the loss of crystal transparency as particles interacting with them create color centers [115]. Fortunately, that damage is self-repairing, but the transparency of the crystal does deteriorate in the short term during each hours-long LHC data-taking cycle and in the medium term during each year of data taking. The regions of the ECAL closest to the beam pipe had lost 96% of their transparency by the time LHC Run 2 was over [117]. The loss of transparency in the barrel was comparatively lower, up to 13%. The evolution of the detector transparency is monitored by a system which injects laser pulses into the crystals.

3.2.4 The hadronic calorimeter

The hadronic calorimeter (HCAL) lies beyond the ECAL. It is divided in four parts: the barrel (HB) and endcap (HE) calorimeters, located within the solenoid, the outer calorimeter (HO), located just outside the solenoid in the central region of CMS, and the forward calorimeter (HF), found close to the beam pipes on either side of the detector. The HCAL is designed to destructively measure the energies of charged and neutral hadrons.

The HCAL is a sampling calorimeter, which means that it is made of alternating layers of absorber and active materials. In the HB, HE and HO, plastic scintillators form the active material, the absorber material varies from one part of the detector to the other and can be brass, stainless steel, iron, or even the solenoid coil. The absorber material triggers hadronic showers, and the particles from these showers excite the scintillators as they go through it, the scintillators then emit photons in the visible spectrum which are collected by photodetectors. The advantage of this sampling setup is that it is possible to combine a material which effectively

triggers hadronic showers with another material with good scintillating properties, instead of having to find a single material capable of fulfilling both requirements. The disadvantage is that part of the shower is lost in the absorber material and never reaches the scintillators. Detectors are calibrated and corrected for this effect, but it adds an extra uncertainty term to the energy resolution.

The absorber material is mostly brass in the HB and HE, but both the innermost and the outermost layers of the HB are made of stainless steel in order not to compromise the detector's structural strength [93]. The HB in the most central regions of CMS is not thick enough to contain all hadronic showers, which is why the HO is located outside the solenoid in the region $|\eta| < 1.3$. Most of the HO consists of a single scintillating layer, using the solenoid coil as an extra absorber layer. However, the most central part of the HO needs even more absorbing material and consists of two scintillating layers on both sides of an absorbing layer of iron.

The hadronic forward calorimeter (HF), which extends the $|\eta|$ coverage until 5.2, operates in a much harsher radiation environment and thus uses a different technology for the active material, based on the emission of Cherenkov light [118]. Cherenkov light is emitted when particles travel faster than the speed of light in the scintillating medium, which in this case is made of quartz fibers. There are both long and short quartz fibers, inserted into a steel structure serving as the absorber material. The fibers are connected at the outer end of the absorber material to light guides connected to photomultipliers. The fibers are parallel to the beampipe, the long ones run all the way through the detector while the short fibers only go halfway. Electrons and photons do not penetrate deeply into the absorber material and thus activate preferentially the long quartz fibers, while hadrons travel deeper into the absorber and thus trigger similar responses from the long and short fibers. This differential response of the fibers allows hadronic and electromagnetic showers to be distinguished from each other.

The photodetectors used in the HB, HE and HO were originally hybride photodiodes (HPDs), but those have progressively been upgraded to silicon photomultipliers (SiPM) over a period of many years. The HO was upgraded before Run 2 [119], the HE in the 2017-18 technical stop, and the HB only after the end of Run 2 [120]. SiPM are a new technology with many advantages over HPDs: they have higher photon detection efficiency and much higher gain. New readout electronics were also installed, allowing a large increase in the number of channels and the depth-segmentation of the detector. The HF was also upgraded, first during Long Shutdown 1 by replacing the photomultipliers tubes (PMTs) with a newer model better able to reject anomalous signals coming from particles directly interacting with the PMTs, and then during the 2016-17 technical stop by replacing the readout electronics and doubling the number of output channels [118].

The HCAL was initially calibrated by results from test beams and by using radioactive sources and lasers, but a substantial improvement comes from using cosmic muon and proton-proton collisions data [121]. The energy resolution in the HB was found in 2002 test beams to be [122]:

$$\left(\frac{\sigma_E}{E}\right)^2 = \left(\frac{115\%}{\sqrt{E(\text{GeV})}}\right)^2 + (5.5\%)^2 \quad (3.8)$$

After calibration using 2016 proton-proton collisions data, the energy resolution in the HB/HE/HO part of the HCAL is about 17.8% for 50 GeV charged pions [121], during normal CMS operations.

3.2.5 The muon system

Gaseous subdetectors dedicated to the tracking of muon trajectories are installed in the outer parts of the CMS detector, within the iron return yoke for the solenoid magnetic field. Of the final state SM particles found in CMS (particles that have a decay time longer than the time it takes them to cross the detector), muons are the second most penetrating: they do not have strong interactions like hadrons, and they cannot convert to a pair of electrons like photons. They can emit bremsstrahlung radiation like electrons, but the emission probability is inversely proportional to the square of the particle mass, which implies that bremsstrahlung for muons is $\sim 1/40000$ as important as for electrons, given the ~ 200 mass ratio between the two particles. Neutrinos are of course more penetrating than muons as they have no electric or color charge and just leave the detector without interacting with it at all. Otherwise, the dense material particles have to cross before reaching the outer part of the detector ensures that no other particles apart from muons interact with the muon system, with very few exceptions⁹. The return magnetic field of the solenoid bends the tracks of muons as they go through the gaseous subdetectors, allowing their momenta to be measured.

There are three types of gas ionization detectors used in the muon system: drift tubes (DTs) in the barrel region, cathode strip chambers (CSCs) in the endcaps, and a complementary trigger system in both regions is provided by resistive plate chambers (RPCs). DTs are located in the region $|\eta| < 1.2$, where background and muon rates are low and the magnetic field is mostly uniform and contained in the iron yoke. In the endcaps, where background and muon rates are high and the magnetic field is non-uniform, the faster response time, finer segmentation and increased radiation resistance of CSCs make them a superior choice, and they cover the pseudorapidity region $0.9 < |\eta| < 2.4$, partially overlapping with the DTs. RPCs have a fast response time to signal and excellent time resolution, while their position resolution is not as good as that of DTs or CSCs. They provide an independent trigger system for $|\eta| < 1.6$, their ability to uniquely identify the specific bunch crossing to which a muon track belongs is particularly useful [93].

The DT system is separated into four stations, one on either side of the iron yoke and two embedded within it, forming concentric cylinders around the inner part of CMS, as shown in figure 3.6. The cylinders are themselves divided into five wheels spread out along the z direction: there is a central wheel around $\eta = 0$ and two wheels on either side of it. Each muon station is composed of many chambers, each containing three superlayers (or two, in the case of the outermost station), and each of these superlayers is made of four layers of cells. The cells are long in one dimension, $\sim 2\text{-}4\text{ m}$, and short in the other two, about 21 mm wide and 11 mm in the radial direction. Two of the superlayers in each chamber are oriented in such a way that they can measure the azimuthal angle ϕ of the muons, while the third superlayer (missing in the outermost station) is located in between and oriented orthogonally to them, and allows the z -position of the muon to be measured. A stainless steel anode wire runs in the center of each drift cell [123] and collects the electrons ejected by the gas atoms ionized by any high energy muon going through the cell. Two of the walls of the cell form the cathodes while the other walls are kept at ground potential. The drift times of the electrons can be measured and used to determine at what distance from the wire the muon crossed the cell, which considerably improves the position resolution. The maximum electron drift time is around 400 ns, which means that a cell can be busy for up to 16 bunch crossings after a signal is detected. This is not a problem due to the low event rate per unit area in the barrel.

The CSCs are trapezoidal, with the short side closer to the beampipe and the long side

9. There are dedicated algorithms to deal with punchthrough from the HCAL, as explained in chapter 5.

further away. They each cover 10° or 20° in ϕ , partially overlapping in order to avoid dead zones. They are organized in four stations within the endcap steel return yoke, as shown in figure 3.7, in such a way that a muon with a pseudorapidity $1.2 < |\eta| < 2.4$ goes through 3-4 CSCs. Muons with smaller $|\eta|$ go through fewer CSCs but this is compensated by them crossing DTs too. The CSCs are multiwire proportional chambers made of alternating planes of anode wires and cathode panels with strips [93]. There are six wire planes and seven cathode panels in each CSC. The cathode strips have constant $\Delta\phi$ and run radially along the trapezoid, while the wires run in the orthogonal direction. The wires are operated in avalanche mode, which means that when a muon goes through the chamber and ionizes the gas, the resulting electrons trigger a chain reaction when they get close to the anode wire: the electric field there is so strong that the drifting electrons acquire enough energy to ionize other atoms, and the resulting electrons ionize other atoms, etc. The avalanche signal is detected by the wire, pinpointing the radial coordinate of the track. The avalanche also leaves near the wire an accumulation of positive ions, which drift slower than electrons, and those ions induce a charge in the cathode strips closest to them, which allows the determination of the ϕ coordinate of the track.

RPCs are present in both the barrel and endcap, and provide a redundant trigger capability in the $|\eta| < 1.6$ region they cover. The basic RPC design consists of two parallel high resistivity plates, separated by a gas gap of a few millimeters [123]. The outer part of each plate is coated in conductive material to form either the high voltage or the ground electrode. Aluminum strips insulated from the electrodes are used to pick up the signal, in a similar process to that of the cathode strips. CMS RPCs are double gap, with the readout strips located in between the two gaps and thus sensitive to discharges from both of the gaps [125]. In order to cope with the high radiation environment of CMS, the RPCs are operated in avalanche mode, with limited gain in the gas volume but high amplification in the front-end electronics, allowing them to handle higher event rates [126].

To achieve a good energy resolution for muons, the locations of the various muon chambers containing DTs, CSCs and RPCs must be known with high accuracy. Some misalignment in the muon system inevitably occurred during the construction of the chambers and the assembly of the detector, but these were measured once the chamber installation was complete. However, misalignment can also come from the magnetic field deforming the return yoke and from time-dependent effects caused by e.g. temperature fluctuations. This is why a continuous monitoring of the absolute and relative positions of the muon chambers is necessary: this is done by an optical alignment system making use of LEDs and laser beams. Light shafts connect the muon chambers to each other and also to the tracker at the center of the CMS detector [93].

The muon momentum resolution is dominated by the measurements of the inner silicon tracker for low p_T muons. As the p_T increases and the tracks get straighter, the resolution decreases. Simulations show that some of the loss in resolution can be recovered by using information from the muon system (cf. figure 3.8). The muon p_T resolution measured in 2016-17 cosmic rays data is shown in figure 3.9. Cosmic ray muons can only travel so far through the Earth, and thus most of them come from directly above CMS, with low values of $|\eta|$, and data is lacking for high p_T high $|\eta|$ cosmic muons. Still, we can observe that for low p_T muons in the 10-100 GeV range, the resolution is at the 1 % level for $|\eta| < 1.2$ and 2 % for $1.2 < |\eta| < 1.6$, but CMS retains a good resolution even for a transverse momentum of 0.6-1 TeV: 4.5 % for $|\eta| < 1.2$ (data is lacking for high p_T $1.2 < |\eta| < 1.6$ muons).

The muon system is scheduled to have an upgrade during LS2 by adding an extra set of detectors in the endcaps: gas electron multipliers (GEMs). In the region $|\eta| < 1.6$ redundancy in the muon system is provided by the presence of RPCs in addition to DTs and CSCs. How-

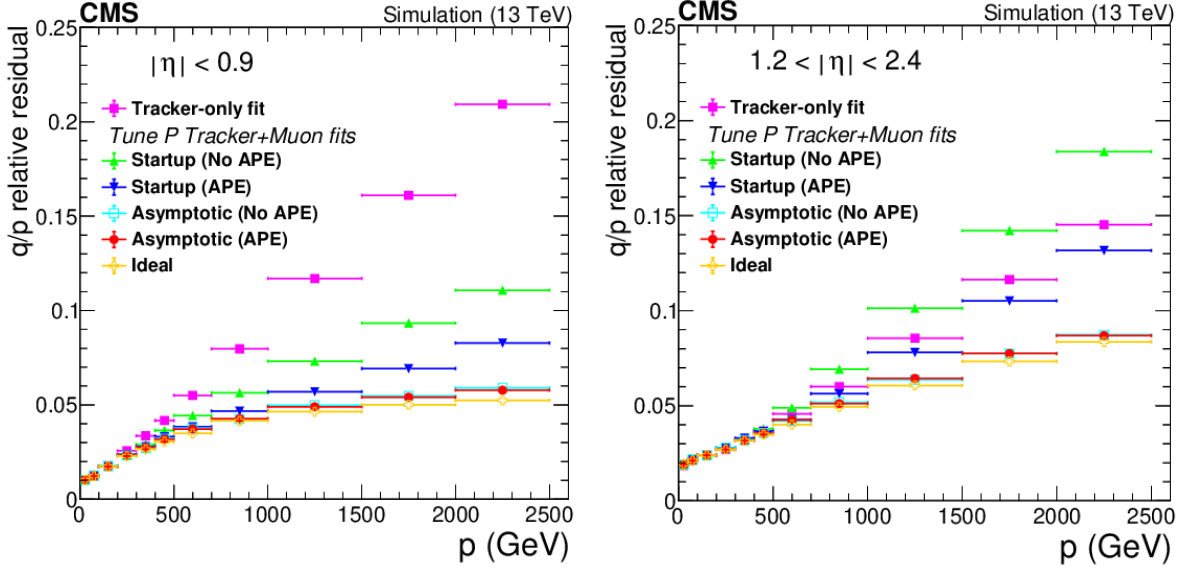


Figure 3.8 – Muon momentum resolution as a function of momentum, estimated in simulations, using only the inner tracker, or using a global algorithm which takes into account information from both the inner tracker and the muon system [127] (TuneP Tracker+Muon fits). Different detector alignment scenarios are shown for the global algorithm. They differ along two dimensions: whether they use the startup or the asymptotic alignment (the startup alignment corresponds to the relative positions of the subcomponents of the detector at the beginning of data-taking while the asymptotic alignment is calibrated with data taken during Run 2), and whether they take into account alignment position errors (APE) or not. The best results are obtained with asymptotic alignment and APE. Left: in the pseudorapidity region $|\eta| < 0.9$, right: $1.2 < |\eta| < 2.4$. Using both the inner tracker and the muon system bring a considerable improvement to the resolution compared to using only the inner tracker, especially for high momentum muons.




Figure 3.9 – Muon p_T resolution as a function of momentum, measured in data and compared to simulations (DY MC) [127]. Left: in the pseudorapidity region $|\eta| < 1.2$, right: $1.2 < |\eta| < 1.6$. There is good agreement between the data and the simulations.

ever, due to the high rates expected after LS2, RPCs are not adequate for the high $|\eta|$ regions of the muon system, and GEMs are a better solution [128]. A first GEM station was installed and commissioned during Run 2 [129], others are being installed as of the writing of this thesis.

Chapter summary

The data analyzed in this thesis comes from proton-proton collisions collected at CMS, with the proton beams provided by the CERN LHC. The LHC is an approximately circular collider operated by CERN and located in the French-Swiss border region, near Geneva. The LHC is fed by a series of other accelerators which shape the bunch structure of the proton beams and increase their energy until they can be injected into the LHC and accelerated further. During Run 2, the data-taking period which is analyzed in this thesis, the instantaneous luminosity delivered by the LHC for proton-proton collisions reached a maximum of $1.9 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, higher than its nominal value, and the collisions were performed at $\sqrt{s} = 13 \text{ TeV}$, the highest center-of-mass energy ever reached by a collider. Both parameters are set to improve in future data-taking runs. CMS is a general purpose detector located at one of the collision points of the LHC, featuring a superconducting solenoid providing within it a homogeneous magnetic field of 3.8 T. A series of subdetectors installed around the collision point provide a variety of measurements and as close to full angular coverage as possible. Within the solenoid, a silicon tracker records the flight path of charged particles and measures their momenta from the way they are curved by the magnetic field, an electromagnetic calorimeter measures by a destructive process the energies of photons and electrons, and a hadronic calorimeter does the same for hadrons. Outside the solenoid, gaseous detectors track the trajectories of charged particles which manage to get through the rest of the detector: these are almost always muons. With this setup, CMS can measure the energies and momenta of all SM final state particles, except for the undetectable neutrinos.

The CMS trigger system

The previous chapter showed how proton beams are accelerated and made to collide, and also described the detector apparatus installed around the collision point. This chapter describes how the trigger system analyzes the raw data generated by particles in the front-end electronics and selects some of it to be kept in permanent storage. Proton-proton collisions happen at a high rate in CMS: the nominal number of bunches in the LHC, 2808, multiplied by the bunch revolution frequency of 11.2 kHz, gives an average collision frequency of 31.4 MHz. However, all of these events cannot possibly be saved in long term storage, the main limitations are the bandwidth of data recording and transfer from CMS to the CERN Tier0 computing site, and the Tier0 capacity for prompt reconstruction of the data. This limits the data recording rate to about 1 kHz [130].

The trigger system is a set of algorithms which accomplish this enormous reduction in the data rate while keeping the most interesting events as defined by the CMS physics program. The system is divided in two levels: a Level-1 (L1) based on custom hardware, explained in section 4.1, which performs a first selection and feeds into a software-based High Level Trigger (HLT), explained in section 4.2, which runs on computer farms and makes the final decision on whether an event is kept in long term storage.

4.1 Level-1 trigger

The Level-1 (L1) trigger is implemented on hardware, on both programmable hardware like field-programmable gate arrays (FPGAs) and on faster, more radiation resistant but less flexible application-specific integrated circuits (ASICs) [93]. A major upgrade was done during Long Shutdown 1 (LS1), with the aim of improving the L1 performance at the higher instantaneous luminosities of LHC Run 2 and simplifying its electronic components, which are now mostly based on a few common hardware designs [131]. The architecture of the L1 after the upgrade is shown in figure 4.1.

The L1 reduces the data rate from the many tens of MHz of proton-proton collisions to a maximum output rate of ~ 100 kHz. The decision on whether to reject an event or send it to the HLT is made within $4\ \mu\text{s}$ of the collision, using information only from the calorimeters and muon chambers [108]. A global trigger makes the final L1 decision, based on both the calorimeter and muon triggers, and also on the readiness of the CMS subdetectors and the data acquisition system (DAQ).

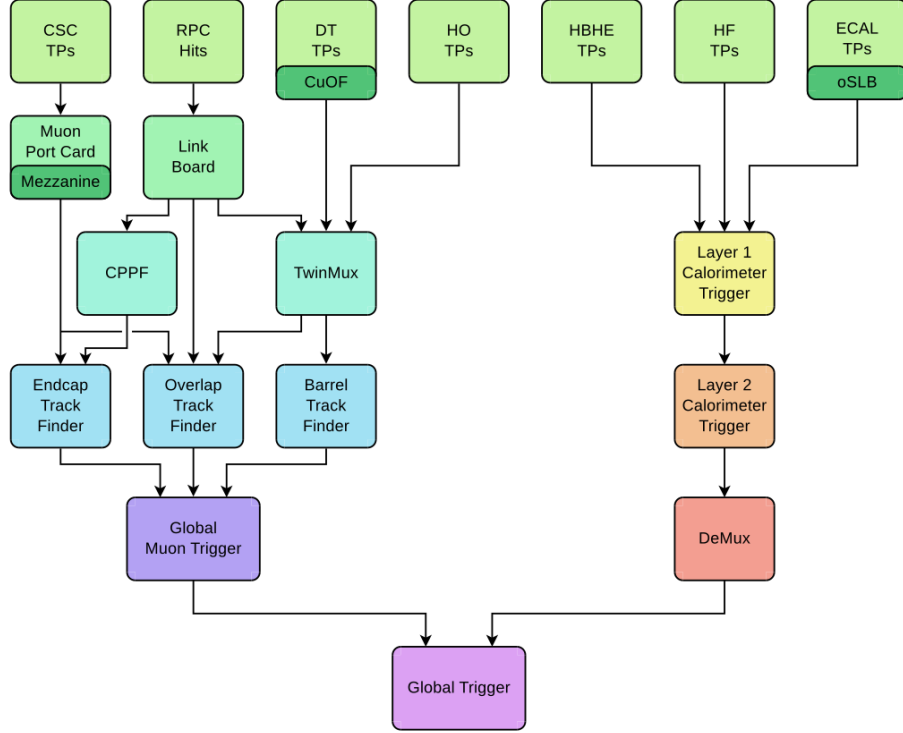


Figure 4.1 – The L1 trigger system after the LS1 upgrade. The components and the relations between them are explained in the text [132].

4.1.1 Calorimeter trigger

The basic data used in the calorimeter trigger comes from energy deposits in the ECAL or HCAL. The ECAL, used to identify electrons and photons, is separated into a barrel (EB) and two endcap (EE) sections. In the barrel, each 5×5 array of PbWO_4 crystals forms a trigger tower, and the transverse energy in each of them is computed and forms a trigger primitive, which is processed in the calorimeter trigger algorithms. In the endcap, trigger primitives are also computed from the transverse energy deposited in trigger towers, but the geometry of these is more complex and the number of crystals per tower varies from 25 to 10 with increasing $|\eta|$ [133]. As mentioned in chapter 3, the ECAL loses transparency when irradiated and recovers some of it when not. This effect is mitigated by using information from the laser monitoring system and regularly updating corrections to the trigger tower energies. The HCAL trigger primitives are also based on transverse energy (E_T) deposits, and can combine many readout channels into a single readout tower in the same way ECAL crystals are combined into trigger towers.

After the LS1 upgrade, the calorimeter trigger began using a time-multiplexed architecture with two layers, Layer-1 and Layer-2. Each electronic card in Layer-1 is mapped to a part of the detector which is four trigger towers wide in ϕ and spans all of η . The Layer-1 cards receive the trigger primitives generated from ECAL and HCAL data in their sector, calibrate and sort them and then transmit them to Layer-2. Time-multiplexing means that for each bunch crossing each Layer-1 card retransmits data to Layer-2 using a single optical fiber link, but this data is sent over many bunch crossings instead of sending all of it immediately [131]. The data from the next bunch crossing is similarly transmitted over many bunch crossings, but to a different Layer-2 card and using a different optical link. After nine bunch crossings, the first optical link has finished transmitting and can be used to transmit data from a new bunch

crossing, and this cycle can be repeated indefinitely.

Layer-2 is set up in such a way that each card receives calibrated trigger primitives from all of the detector for a single bunch crossing. Due to the time-multiplexing, they receive and process this data progressively, performing a one-dimensional scan along the η direction¹ of the detector, while constructing and sorting physics objects such as electrons/photons, jets², taus and calculating global sums such as missing E_T [134]. A final demultiplexing stage prepares the data to be sent to the global trigger, which makes the L1 Accept decision.

4.1.2 Muon trigger

All three types of gaseous detectors located in the muon chambers contribute data to the muon trigger, which is separated according to pseudo-rapidity regions: there are barrel, overlap and endcap muon track finders, respectively abbreviated BMTF, OMTF and EMTF. The trigger primitives constructed by DTs and CSCs are local track segments, while the RPCs provide a collection of hits.

The DTs are equipped with custom and programmable electronics which reconstruct up to four local track segments per bunch crossing per chamber [131]. The local track segments are assumed to point to the interaction vertex and include information on their radial position, their bending angle, and the number of layers used to reconstruct them. Despite the 400 ns maximum electron drift time (about 16 LHC bunch crossings, cf. chapter 3), the DT trigger electronics are able to identify the correct bunch crossing with an efficiency of 97-99%, except for tracks with a high angle of incidence, for which it drops to $\sim 88\%$ [135]. The local track segments coming from all DTs are synchronized and transmitted to the BMTF and the OMTF (via the TwinMux, described later in this section).

The CSC electronics reconstruct separate track segments from anode and cathode hits, these track segments are then correlated. Each anode track segment has a well-defined radial distance from the beam line as well as precise timing information, the corresponding cathode track segment provides the ϕ coordinate. Up to two local charged track segments per bunch crossing are sent by each CSC detector to an electronic component called muon port card [131], which forwards the track segments to the EMTF and OMTF.

DT track segments and RPC hits are not directly fed to the track finders, they go first through another module, either the TwinMux or the CPPF. The TwinMux module has a dual role [131, 136]:

- Combining DT track segments and RPC hits into “superprimitives” to be sent to the BMTF. These superprimitives have both the high spatial accuracy of the DTs and the high timing accuracy of the RPCs.
- Consolidating the data coming from the DTs and RPCs in many slow links into fewer high speed output links connected to the BMTF and OMTF. The input data coming from a single sector is also duplicated and sent to the track finder processors working on neighboring sectors, thus avoiding the need for information sharing between the processors.

The CPPF performs a similar data consolidation and duplication job for RPC hits to be sent to the EMTF.

1. This one-dimensional (rather than two-dimensional) scan of the detector allows for a great simplification of connections between different electronic components, greatly improving the data rate which can be analyzed.

2. Jets are collections of tightly packed particles, mostly made up of hadrons. See chapter 5 for more information.

The muon system barrel, where the BMTF operates, is divided into wheels and wedges. There are 12 wedges with a trapezoidal shape, they have a limited $\Delta\phi$ spread but run along the entire z axis of the barrel. As mentioned in chapter 3, the barrel is divided into five wheels along the z axis. The intersection of a wedge and a wheel corresponds to a sector, and each of them has a dedicated track extrapolation unit, which also receives duplicated data from the nearest neighboring sectors. The track extrapolation units check if a superprimitive from an outer layer is compatible with a superprimitive from an inner layer, by making use of a look-up table which takes as inputs the inner superprimitive coordinates and bending angle and then determines an acceptable ϕ window for the outer layer. If the outer layer superprimitive is within that window, the two superprimitives are paired up since they are considered compatible with the same muon track. A track assembler unit takes these superprimitive pairs created in all layers of the detector and constructs complete tracks by combining them. A quality is assigned to each track according to the number of layers which contribute to it.

The OMTF operates in the barrel/endcap overlap region and thus receives information from the DTs, RPCs and CSCs. The algorithm reconstructs tracks independently in each ϕ sector, starting from reference hits. Up to four different reference hits are selected, favoring those from inner layers and with better ϕ resolution. A reference hit from a single layer is associated to hits in other layers according to whether they match any of 26 different “golden patterns” corresponding to different muon transverse momenta. When multiple golden patterns are compatible with the reference hit, the one matching the highest number of hits (i.e. the highest quality one) is preferred.

The EMTF uses data from the CSCs and the RPCs located in the endcaps. The algorithm reconstructs tracks using at most one trigger primitive (CSC or RPC) per layer: by default the better resolution CSC is preferred, but the RPC can be used if the CSC system missed a hit in one layer. The first step is looking for correlations in ϕ between different layers, checking if they correspond to any of five predefined patterns, this is done in parallel in four η sectors. The resulting tracks are ranked according to quality criteria determining how precisely their p_T can be measured: their straightness, the number of layers used to construct them, and whether the two inner stations (where the magnetic field is higher) are involved. The three highest quality tracks are saved and have their p_T 's evaluated by a machine learning algorithm³ implemented in a look-up table. To evaluate the p_T of a muon, one would normally use ϕ bending angles in each layer and ϕ differences between different layers. However in the endcaps there are multiple factors complicating this evaluation: the magnetic field is non-uniform, high p_T muons may trigger electromagnetic showers when interacting with the detector material while low p_T muons may be scattered off their trajectories or lose large amounts of energy. A machine learning algorithm trained on large amounts of simulated data can make sense of this complicated picture. Its inputs and outputs are implemented in a look-up table so that it can make fast p_T assignments [137].

The upgraded global muon trigger (μGMT ⁴) receives the best muon candidates from each of the three track finders. The μGMT ranks the muons by p_T and quality and removes duplicates which may be produced by adjacent sectors or wedges in the same track finders, or muons which are reconstructed both by the BMTF and the OMTF, or both the OMTF and the EMTF. In parallel, the μGMT also extrapolates the muon track coordinates back to the

3. More precisely, a Boosted Decision Tree (BDT), explained in more detail later in chapter 5.

4. It is called upgraded global muon trigger as opposed to the previous global muon trigger used in LHC Run 1.

collision vertex via a look-up table; this allows e.g. better reconstruction of the invariant mass⁵ of multiple-muon systems. The eight best tracks, with their extrapolated coordinates, are sent to the upgraded global trigger (μ GT) for further processing. Quality information is also transmitted: muons from the BMTF are always considered to be of the highest (“tight”) quality thanks to the strong magnetic field in the barrel, while muons from the OMTF and EMTF are classified into either “tight”, “medium” or “low” categories according to how many layers they hit and whether they hit the first layer.

4.1.3 Global trigger

The upgraded global trigger (μ GT) makes the final L1-Accept decision based on information from Layer-2 of the calorimeter trigger and the μ GMT. The incoming data consists of global sums, such as the missing transverse energy, and of candidate particles (muons, isolated and non-isolated electrons/photons, jets, taus) with information about their E_T or p_T , η and ϕ coordinates, and quality of the reconstruction. There is also information coming from CMS subsystems which can be used to accept or veto events in standalone “technical triggers”, or used in combination with calorimeter or μ GMT data for a trigger decision. Technical triggers are useful for the calibration and commissioning of the various CMS subsystems, including the L1 trigger itself [108].

L1 algorithms can select events based on several different criteria, e.g. if a given particle passes selections on their p_T/E_T , η or ϕ , or if two different particles are spatially correlated, or if a given global sum is higher than a certain threshold. Different criteria can also be combined with logical operations to produce more complicated triggers. After passing all of these criteria, an event may still be subjected to a so-called prescale before being accepted by the L1 trigger. Prescales are positive integers, including zero, determining how many events fulfilling all other requirements of a specific L1 algorithm may actually fire the trigger. A prescale of zero means that no event will be accepted, this is typically used for calibration triggers when the instantaneous luminosity is high and the priority is to collect data useful for physics analyses. Otherwise, applying a prescale $N \geq 1$ means that only one out of N events passing all other trigger requirements will be accepted, effectively dividing the trigger rate by N .

The collection of L1 algorithms encoded in the μ GT firmware is called the L1 trigger menu, and since it is fully implemented in programmable electronics, it can be changed when data-taking conditions change. Up to 512 different L1 algorithms can be used, and in practice 350-400 were used during Run 2. For hardware reasons, the L1-Accept rate cannot exceed 100 kHz, but it is otherwise desirable to keep it as high as possible. So over the years, as the instantaneous luminosity and the center-of-mass energy increased, new L1 menus with reduced rates had to be developed. Different strategies can and were used to reduce trigger rates:

- Increasing the prescale, e.g. from 2 to 10, which in this case would mean a new rate equal to 1/5 of the old rate. This solution is however not desirable for triggers used in physics analysis for which one would like to keep as much data as possible (i.e. the desirable prescale is 1). Prescaled triggers are fine for calibration or data validation studies.
- Tightening the requirements on the physics object, e.g. by requiring an electron/photon candidate to be isolated (i.e. requiring little activity in a small region of the detector around the candidate), or restricted to a certain $|\eta|$ range.
- Increasing the p_T/E_T thresholds. These can however be set too high for some specific analyses, e.g. a 50 GeV threshold on the muon p_T severely reduces the rate of $Z \rightarrow \mu\mu$ events triggered.

5. See appendix A for a definition of the invariant mass.

- Using multiple-object instead of single-object triggers. If e.g. two muons are required instead of one, the p_T thresholds on both of them can be kept lower.

The data analysis presented in this thesis considers only high energy events, which means that single-object triggers with high p_T/E_T thresholds can be used with no risk of losing signal efficiency.

The instantaneous luminosity not only varies over the years of LHC operations but also during a single LHC fill⁶. When the machine has just filled up with two stable beams and they start colliding, the instantaneous luminosity is at its highest and then starts decaying, as explained in chapter 3. In order to keep the rates high and not waste any data, different sets of prescales, called “prescale columns”, are used. For lower luminosities, lower values of the prescales are used. The CMS trigger shifter is in charge of changing the prescale column when the luminosity gets below a certain value. They also monitor the L1 rates (in particular checking that the total rate is less than ~ 100 kHz) and the various L1 subsystems⁷. The μ GT provides a preview of the rates of other prescale columns, making it easy to check that the new rate will stay below the 100 kHz limit.

4.2 High Level Trigger

The High Level Trigger takes events for which there was a L1-Accept signal and decides if they should be kept for permanent storage, reducing the rate from ~ 100 kHz to an average of ~ 1 kHz, a number limited by the data transfer bandwidth and the CERN Tier0 computer center capacity for prompt data reconstruction. This large rate reduction is performed by a collection of ~ 600 algorithms, also called HLT paths, which collectively form an HLT menu. Unlike the L1 trigger, the HLT is implemented in software running in a computer farm.

4.2.1 HLT algorithms

An HLT path only starts running if one of its L1 preconditions is fulfilled. Then a series of object producers and filters are applied. Producers can build objects as simple as ECAL energy deposits or as complicated as global muon tracks combining inner tracker and muon chamber information. Filters check whether the objects produced fulfill the trigger requirements and reject the rest. Producers and filters can be shared between different HLT paths in order to save computing resources. Another way to save resources is to make sure producers constructing simple objects are run before producers reconstructing tracks or particle candidates using the Particle Flow (PF) algorithm⁸, this way many events are rejected before any complicated algorithms need to be run [138]. Figure 4.2 (left) shows the distribution of HLT decision times in a specific 2016 data-taking run. The distribution is heavily skewed, peaking under 0.1 s, but some HLT decisions can last several seconds. The right-hand side of the figure shows the increase of the average processing time as the instantaneous luminosity increases. The higher luminosity runs of 2016 got close to going over the HLT timing budget, which was around 220 ms at the time. More CPU cores were purchased for the 2017-18 data taking in order to cope with the increased luminosities, allowing a maximum processing time of 320 ms in 2018, which in practice can be increased by 20% thanks to hyperthreading [130].

6. A fill is a data-taking period lasting as long as two specific beams are circulating in the LHC. Dumping the beams marks the end of a fill, a new fill starts when new beams are injected into the machine.

7. I did about two weeks of trigger shifts at CMS Point 5 during 2018 data-taking.

8. PF is explained in more detail in chapter 5.

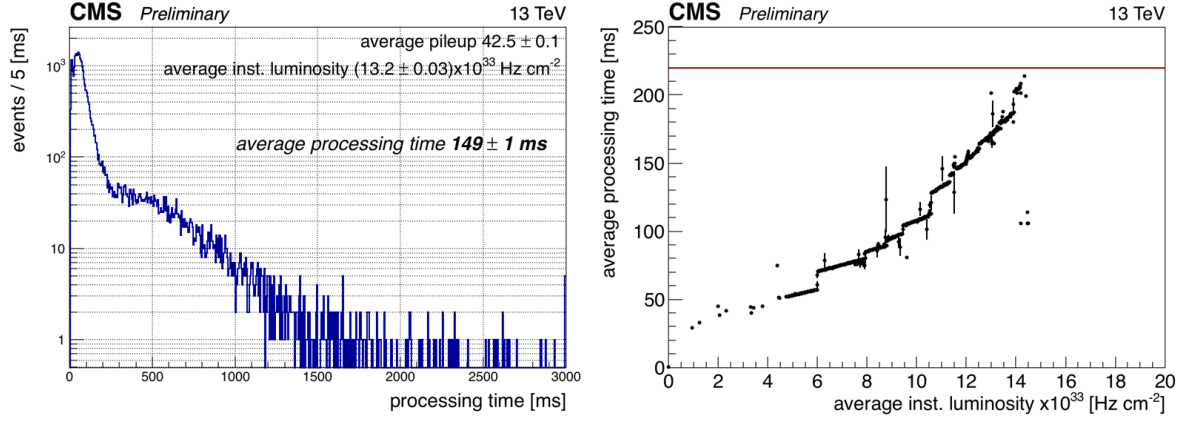


Figure 4.2 – (Left) Distribution of HLT processing times for a 2016 data run when the average instantaneous luminosity was $\sim 13.2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ and the average number of simultaneous events (pileup) was 42.5. (Right) Average HLT processing time as a function of instantaneous luminosity. The processing time increases faster than linearly when instantaneous luminosity increases, the red line shows the maximum average processing time that the HLT computer farm as it existed in 2016 could handle. Later upgrades pushed this limit to higher values [138].

An event passing all filters of an HLT algorithm may still be subject to prescales before being accepted for permanent storage. Like for the L1 trigger, many different prescale columns are defined for different luminosities. Accepted events are written into separate datasets which group together the outputs of many separate HLT paths, e.g. there is a SingleMuon dataset for events containing a high momentum muon. The datasets are non-exclusive since the same event can fulfill the requirements of two different HLT paths belonging to different datasets. There are of course efforts to ensure that the overlap between different datasets is minimal as duplicated events waste computing resources, and in fact, as detailed later in this section, I personally made a study to check the event rate overlap between different datasets and recommended the merging of those which shared a considerable number of events.

The 1 kHz limit on the average HLT rate can be bypassed in two ways:

- By having a much reduced event size. This is achieved in “scouting” data, by saving only objects reconstructed at HLT [139]. The same computing resources allow the collection of many more events.
- By “parking” the data, and reconstructing it later when computing resources are available rather than reconstructing it immediately after it was collected. Special parking triggers for studies of b-quark physics were used in 2018 [130].

Figure 4.3 shows how the HLT rate evolved over the duration of a specific 2018 LHC fill. The promptly-reconstructed “physics” data rate starts at about 1300 Hz and finishes at around 600 Hz, for an average close to the 1000 Hz limit over the whole fill. In comparison, the parked data has a much higher rate. There are also calibration and commissioning triggers, the rates of which are not shown here.

4.2.2 HLT rates monitoring

HLT paths are not centrally designed, they are developed by physics analyzers who are better placed to know what trigger would be best suited for their analysis. There is however a Trigger Studies Group (TSG) ensuring the coordination of trigger development, and the deployment of new HLT menus. The performance of said menus is also checked, in particular

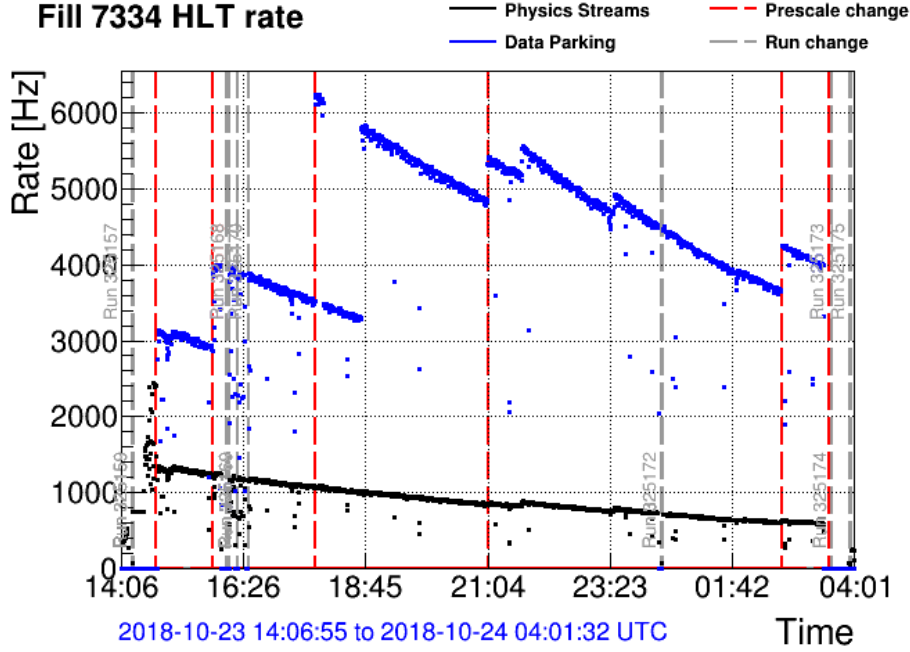


Figure 4.3 – HLT rates for LHC fill 7334 (2018, taken from the CMS trigger rate monitoring website). Rates for the promptly-reconstructed “physics” data streams are shown in black, the much higher rates for the parking streams are shown in blue. The vertical red dotted lines indicate changes in prescale columns. The grey dotted lines show the start of new runs. Runs are technical subdivisions of CMS data, a rapid succession of new runs usually indicates some technical problem.

whether the rates and processing times stay within reasonable limits. Another consideration is that CMS data is useful for many different kinds of analyses, such as H boson physics, searches for supersymmetry, or searches for new high mass objects, which is why the share of trigger rates each of these analysis groups gets is monitored to check if it corresponds to CMS policy. There are also different object groups which check the performance of the reconstruction and identification of specific objects such as muons or taus, and the rates of each of these groups also need to be monitored. For this purpose, the plots shown in figure 4.4 can be useful, I produced them for the trigger monitoring subgroup (STEAM) of TSG.

Over the course of my thesis, I worked at STEAM and estimated the rates of many Run 2 HLT menus before they were deployed. At the beginning I used a legacy code before working with other STEAM colleagues to develop a new leaner code with a much faster run time. I eventually became the main person responsible for maintaining and updating the code. It can be used to estimate rates of individual HLT paths, the rates of datasets and how they overlap, and the rates assigned to different physics analysis and object groups as already shown in figure 4.4. All of that can be done using either data or simulations, each of which have their own advantages and drawbacks. The data used consists of events passing the L1 selections with no HLT requirements other than a large prescale in order to keep the data collection rate low enough. The data is most useful when a new HLT menu is deployed with no changes to the CMS detector or the L1 menu, in this case it is simple enough to emulate the effects of the new HLT paths over existing L1-Accept data. However, when the L1 menu or the detector conditions change too much, data no longer gives accurate estimations and the use of simulations can be preferable, although these have their own accuracy limitations, especially when it comes to wrongly-identified particles or the estimation of the missing transverse energy.




Figure 4.4 – Rates for the different physics analysis and object groups of CMS in one of the 2017 HLT menus. On the left, the rates of each individual group are shown without considering overlaps. On the right, overlaps are considered and events useful to N different groups are counted as $1/N$ for each of the relevant groups. The groups are shown by their abbreviations. Among the physics analysis groups one can find the H boson group (HIG), the supersymmetry group (SUS), and the group searching for new high mass objects (EXO for “exotica”). Among the object groups one can find the tau (TAU), the muon (MUO) and electron-photon (electron-gamma, EGM) object groups.

The rates estimation code works in the following way:

- First the HLT menu is applied to the data or the simulations. The saved output files contain only the trigger information: for each event and each trigger path, a boolean indicates whether the trigger fires or not. This first step is usually performed by subdividing it into many parallel jobs submitted to a computer cluster, in order to quickly process large numbers of files.
- Then, a trigger counting code counts how many times each individual trigger path fires, but also how many times each dataset registers an event, or how many times a trigger corresponding to a specific physics object group fires. When an event fires triggers belonging to different datasets, this overlap is also counted. There is a file summarizing to which dataset and which group(s) each trigger path belongs. This file needs to be updated when new HLT menus are developed. The trigger counting code also keeps track of the number of times any physics trigger fires, without double counting when more than one trigger fires in a single event. Physics triggers are defined as those which are useful to data analyses, either to look for signal events, or as control triggers to e.g. check a particle identification efficiency. The counting task is also usually subdivided into many parallel jobs.
- Finally, a script merges the outputs of the counting jobs, and normalizes the counts in order to obtain trigger rates in Hz. For data, the normalization involves multiplying the trigger count of the HLT path $N_{\text{path}}^{\text{data}}$ by the prescale N_{PS} used to record the data, and dividing the number obtained by the length of time during which the data was taken. CMS data is recorded in indivisible chunks called “luminosity sections” (LS) which last $T_{\text{LS}} = 23.31\text{ s}$, so the duration of data-taking is simply the number of LS N_{LS} times T_{LS} . The rate thus obtained can then be linearly scaled if one wishes to quote the rate for a reference luminosity which is different from the luminosity of the input data. The normalization equation for data is:

$$R_{\text{path}}^{\text{data}} = \frac{\text{reference lumi}}{\text{input lumi}} \cdot \frac{N_{\text{path}}^{\text{data}} N_{\text{PS}}}{N_{\text{LS}} T_{\text{LS}}} \quad (4.1)$$

This linear luminosity scaling is not generally valid: some HLT paths have a non-linear dependence on luminosity. However, it works reasonably well for the small extrapolations performed with this code, typically $\sim 2.0/1.8$.

For simulations, the total rate at which a given process happens in CMS is obtained by multiplying its cross section σ by the reference instantaneous luminosity. However, not all events fire the trigger, so this number has to be multiplied by the share of simulated events firing the trigger: $\frac{N_{\text{path}}^{\text{sim}}}{N_{\text{total}}^{\text{sim}}}$. The normalization equation for simulations is:

$$R_{\text{path}}^{\text{sim}} = \sigma \cdot \text{reference lumi} \cdot \frac{N_{\text{path}}^{\text{sim}}}{N_{\text{total}}^{\text{sim}}} \quad (4.2)$$

4.2.3 Dataset overlap study

I made a study on the overlap of the different physics datasets as they existed at the end of 2017. As mentioned before, CMS events accepted by the HLT are written into non-exclusive datasets, which can lead to duplicated events. There are procedures for removing such events for analyses which use more than a single dataset, the reason these events are undesirable is that they waste precious bandwidth and prompt reconstruction resources which could be better used on other, unique events. It was not an option to simply merge all datasets into one due to computing limitations, but since the division of datasets had not been reconsidered since they had been created at the beginning of Run 1, some advantageous changes were likely to be found.

The study focuses on so-called “physics” datasets, which are useful to CMS data analyses. Here is a complete list of these datasets as they were organized in the beginning of Run 2 (2015-17):

- **BTagCSV**: This dataset gathers trigger paths requiring the presence of two or more jets, with at least one of them being b-tagged by the combined secondary vertex (CSV) algorithm. The CSV algorithm identifies jets which likely originate from the creation of a b quark, as explained in more detail in chapter 5.
- **BTagMu**: These triggers are used to measure the CSV b-tag performance in data. An independent b-tagger is used, one that relies on the presence of a well-identified muon within the cone of a jet. This event signature is associated with the decay of a heavy-flavor hadron⁹, because light-flavored hadrons are either too long-lived to decay within the detector, or they have different decay signatures (e.g. neutral pions decay almost immediately to a diphoton pair).
- **Charmonium**: These triggers are used to study the J/ψ resonance. They select two low p_T muons compatible with the resonance mass of about 3 GeV. Sometimes other requirements such as pseudo-rapidity (η) restrictions are added.
- **DisplacedJet**: Events with displaced jets and other displaced tracks are selected. Displaced tracks are not compatible with the primary interaction vertex. Another selection which is often used is a minimum H_T requirement. H_T is the total scalar transverse energy sum of all jets [132]. These triggers are useful mostly for the search of new long-lived particles. Massive long-lived particles may stop somewhere and later decay, far from the primary interaction vertex.
- **DoubleEG**: Triggers selecting events with two electrons or photons, as well as cross-triggers requiring one electron and a jet.
- **DoubleMuon**: These are typically events with at least two muons.

9. A hadron containing a heavy flavor quark, i.e. a quark which is not up, down or strange.

- **DoubleMuonLowMass:** Events with at least two muons, compatible with the decay of a low mass object.
- **HTMHT:** The trigger paths mostly select events based on H_T , missing H_T , or MET requirements.
- **JetHT:** Events are selected based on a mix of jet p_T threshold and H_T requirements. Triggers selecting “fat” jets¹⁰ are also included.
- **MET:** The trigger paths here select events with large missing transverse energy (MET). Standalone triggers with no other requirements have a MET selection of at least 200 GeV, though thresholds can go lower if filters are applied to remove events with fake MET. There are also cross-triggers with other requirements such as the presence of muons or isolated tracks. The minimum MET thresholds on these cross-triggers can be much lower since the trigger rates are drastically reduced by the other requirements.
- **MuOnia:** The goal is to select two low p_T muons coming from the decay of an Υ meson.
- **MuonEG:** Events with at least one muon and one electron or photon are selected. Sometimes an extra requirement, such as the presence of a second muon, is added.
- **NoBPTX:** These triggers select events in which there was no collision in the bunch crossing. Like the DisplacedJet dataset, this dataset is mostly useful for searching new massive long-lived particles, which may stop in the detector and decay long after the collision. Events with no bunch-crossing collision have far less detector activity than regular events, facilitating the search for long-lived particles.
- **SingleElectron:** The presence of a single high p_T electron is required. Different algorithms may be used to identify the electron. Cross-triggers requiring also the presence of jets or taus are included as well.
- **SingleMuon:** The presence of a single high p_T muon is required. The muon may or may not be required to be isolated. Cross-triggers with tau, jet or H_T selections are included.
- **SinglePhoton:** Events with a single high p_T photon are selected. The simplest triggers include only a p_T threshold, but others have some quality requirements as well. Some paths also include MET selections.
- **Tau:** Triggers requiring the presence of one or more taus are collected here. Some cross-triggers requiring the presence of one muon and one tau are also included. These can in practice function as ditau triggers, since a tau lepton can decay into a muon and two neutrinos.

Using the rates estimation code previously described, I produced the dataset overlap plot shown in figure 4.5, with the datasets as they existed in 2017. The HLT menu used for the study was deployed for a few days of proton-proton collisions data-taking in November 2017¹¹. I used data collected in 2017 to perform the study, rather than simulations. Three groups of datasets showed considerable overlap between them:

- JetHT-BTagCSV-HTMHT, especially the large JetHT-BTagCSV overlap, which corresponds to 1/3 of the total BTagCSV rate and pops out immediately in the graph. This large overlap can be explained by the fact that these datasets all collect different kinds of jet trigger paths.
- SinglePhoton-SingleElectron-DoubleEG. Electrons and photons are similar objects which leave large energy deposits in the ECAL.

10. AK8 jets, see chapter 5.

11. More precisely, it was the 2017 v4.1.0/V6 HLT menu. The v4.2.1 menus were used for the final few days of proton-proton collisions.

- SingleMuon-DoubleMuon. However, SingleMuon already has one of the highest rates out of all CMS datasets, and since the computing limits are not precisely known, it was thought imprudent to make this dataset larger.

There are some comparably-sized overlaps including the missing E_T (MET) dataset, but its rate was already planned to be reduced in future HLT menus.

I then checked how the dataset rates and their overlaps would evolve after doing the two advantageous merges I had identified: JetHT-BTagCSV-STMHT and SinglePhoton-SingleElectron-DoubleEG. The results are shown in figure 4.6. The new datasets, with the placeholder names JetSTMHTBTagCSV and EG, both had lower rates than the SingleMuon dataset and would thus cause no problems to the CMS prompt data reconstruction. Before the dataset merges, the overlap corresponded to around 12% of the total physics rate, which counts only unique events useful for physics analysis or physics object groups. The merges reduce the overlap rate by 38%, an amount corresponding to around 4.6% ($= 0.38 \times 12\%$) of the total physics rate. This means that an extra 4.6% of unique data events could be collected by CMS without changing any of the trigger algorithms or prescales. I presented these results to the Physics Performance and Dataset group, which followed my recommendations and did the two merges, resulting in new datasets called JetHT and EGamma, used during the 2018 data-taking.

Dataset-Dataset Overlap Rates (Hz)




Figure 4.5 – Dataset rates and their overlaps, using the 2017 v4.1.0/V6 HLT menu and the physics datasets as they existed at the time. The plot is symmetric with respect to the diagonal, where the total rates of specific datasets can be found. The rates were estimated using data and scaled linearly to correspond to data-taking at an instantaneous luminosity of $1.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

New Dataset-Dataset Overlap Rates (Hz)




Figure 4.6 – Dataset rates and their overlaps after performing the two merges JetHT-BTagCSV- HTMHT and SinglePhoton-SingleElectron-DoubleEG, the parameters are otherwise unchanged with respect to figure 4.5. The new datasets have the placeholder names JetHTMHTBTagCSV and EG, all other datasets are unchanged and keep their old names.

Chapter summary

The data collected by CMS cannot all be kept in permanent storage. If the LHC is running under nominal operating conditions, the collision rate in CMS and ATLAS is 31.4 MHz, while computing constraints limit the CMS data recording rate to about 1 kHz when averaged over an LHC fill. The trigger system ensures that this rate reduction can happen while keeping the most interesting events in long term storage and rejecting the rest. The CMS trigger is split into two stages: the Level-1 (L1) trigger, implemented on custom and programmable electronics, and the High Level Trigger (HLT), implemented as software in computing farms. The L1 trigger reduces the event rate to 100 kHz, and is split into calorimeter and muon triggers, with a global trigger making the L1-Accept decision. The HLT takes as input events accepted by the L1 and makes the final decision on whether they are kept into long term storage or not. Events accepted by the HLT are written into partially overlapping datasets. The overall HLT rates and the average processing time are monitored by STEAM. A study was performed to check how much overlap there was between different datasets and how it could be minimized, it resulted in the merger of some of them for the 2018 data-taking.

My personal contributions

I worked for STEAM since 2016 at the beginning of my thesis, performing the following tasks (sometimes in collaboration with other people):

- Estimating the HLT rates for trigger menus before they were deployed for data-taking. I estimated rates for individual HLT paths as well as the overall HLT rate and the rates for datasets, physics analysis groups and physics object groups.
- Developing and maintaining a new, lighter framework for the HLT rates estimation, allowing the work to be completed in a single day rather than several days.
- Studying the overlap between the different CMS datasets. I identified groups of datasets which had large rate overlaps with each other and recommended the Physics Performance and Dataset group to merge them and gain $\sim 4.6\%$ in data collection efficiency. They followed my advice, changing the datasets for the 2018 data-taking.

I also did some trigger shifts in the CMS control room during the 2018 data-taking. During each eight-hour shift, I monitored the trigger system in real time, calling experts when there was a problem, and changing the prescale column when the luminosity went below predetermined thresholds.

Object reconstruction in CMS

The raw data collected by CMS needs to be reconstructed in a way that can be used by physics analyses. The momenta and positions of final state particles need to be determined, the nature of these particles needs to be identified (are they muons, electrons, hadrons?), and global quantities such as the missing transverse energy need to be calculated. Some amount of reconstruction is already performed by the HLT in order to select events, and more refined but closely equivalent algorithms are used offline to better reconstruct and identify particles.

This chapter starts in section 5.1 with a general description of the reconstruction of particle tracks, interaction vertices and energy clusters in the calorimeter. The following sections explain how various objects relevant to this thesis are reconstructed, starting with muons in section 5.2, then electrons in section 5.3, photons and hadrons in section 5.4, jets in section 5.5, and the missing transverse energy in section 5.6. The chapter closes with a description of tau reconstruction and identification in section 5.7, and a detailed discussion of a study that I performed about the tau identification efficiency.

5.1 Reconstruction of tracks, vertices and calorimeter clusters

The baseline object reconstruction in CMS is done by a set of algorithms called Particle Flow (PF), which correlates data from all the subdetectors in order to identify final state particles and classify them into non-overlapping categories: muons, electrons, photons, and charged and neutral hadrons [140]. Different particles leave different detector signatures, as illustrated in figure 5.1. The correlation between tracker and calorimeter information allows for a more precise determination of the particle coordinates, energies and transverse momenta. This data can also then be used to construct global sums such as H_T (the total scalar transverse energy sum of all jets) or the missing E_T . Some objects, especially those with high energy, may be reconstructed outside of the PF framework as this can bring about an increase in efficiency. This is the case for electrons and muons in the high mass LFV analysis. However, as explained in section 5.6, some care needs to be taken when combining PF objects with non-PF objects.

Charged particles have an associated track, which for muons may contain both hits in the inner silicon tracker and in the outer muon chamber. There are many different kinds of tracks in CMS, depending on which particle is going through the detector, on the environment of that particle, but also on random factors which affect track quality (some hits may be missing). Different reconstruction algorithms are optimized for different kinds of tracks, and CMS thus uses an iterative tracking algorithm, performing first the track reconstructions which demand

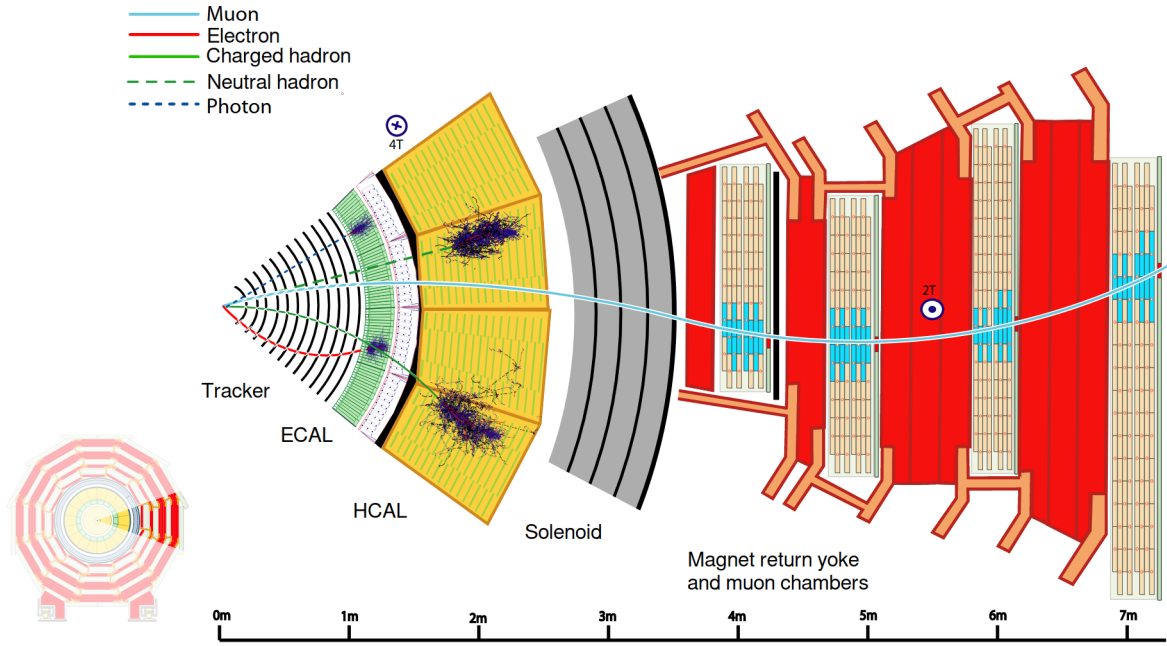


Figure 5.1 – Transverse view of the CMS detector, from left to right: the silicon tracker, the ECAL, the HCAL, the solenoid magnet, and the muon system. Different particle signatures are shown. Muons leave hits in both the inner tracker and the muon system, while leaving little to no energy in the calorimeters. Electrons leave hits in the tracker and then shower in the ECAL. Charged hadrons have reconstructed tracks and reach the HCAL, where they shower. Neutral hadrons also shower in the HCAL but leave no hit in the tracker. Photons shower in the ECAL. They are neutral and thus tend not to leave hits in the tracker, but this can be complicated by conversions to electron pairs. Figure adapted from [141], with labels changed to English.

the least computing resources, removing the hits of any successfully reconstructed tracks, and then moving on to more demanding algorithms over the next iterations. Removing the hits corresponding to already reconstructed tracks is a crucial step that simplifies the combinatorial complexity of further reconstructions. Each iteration of the track reconstruction algorithm proceeds in the following way [142]:

- Tracks are seeded by two or three hits, determining an initial estimate of the track parameters and their uncertainties.
- The seeded trajectories are extrapolated using a Kalman filter [143], taking into account the effects of the magnetic field and possible multiple scattering interactions or other interactions with the detector. Hits compatible with the extrapolation are added to the track candidate.
- After a track is found, a fit of its parameters is performed with a Kalman filter and smoother.
- Tracks are only kept if they pass certain quality flags.

The main way the many steps of the iterative algorithm differ from each other is in the seeding criteria and the quality flags. There are ten well-defined iteration steps, with the first one looking for high p_T tracks originating from the interaction region, also known as high p_T prompt tracks. The next step looks for high p_T tracks originating from secondary vertices (mostly from B hadron decays), then further steps move on to low p_T prompt tracks, lower quality high p_T prompt tracks, progressively more displaced tracks, and individual tracks within high p_T jets. The two final steps reconstruct muon tracks, extrapolating first from the inner tracker out then from the outer muon system in [140]. Track reconstruction for electrons is different because of their high probability of emitting a bremsstrahlung photon while traversing the silicon tracker, this is explained in detail in section 5.3.

After the tracks have been reconstructed, it is possible to locate all proton-proton interaction vertices. This is done by a deterministic annealing algorithm inspired by statistical physics, which finds the global minimum of a problem with many degrees of freedom in a similar way to a physical system approaching its state of minimum energy by gradually reducing its temperature [142]. The algorithm uses as inputs tracks which likely originate from the primary interaction region, which are identified by imposing requirements on the number of pixel and strip hits, on the track fit χ^2 , and on their transverse impact parameter with respect to the beam spot¹. At the start of the deterministic annealing algorithm, the “temperature” parameter is high and all input tracks are compatible with a single proton-proton vertex. The temperature is then progressively decreased, with the algorithm sometimes favoring a splitting of one vertex into two nearby vertices. If the temperature went all the way down to zero, no two tracks would be assigned to the same vertex, there would be as many vertices as there are tracks. The procedure must stop before reaching that point, in a compromise between avoiding accidentally merging two distinct nearby vertices and avoiding creating a spurious extra vertex. When candidate vertices are identified another algorithm performs a final fit on the vertex parameters, determining the three-dimensional positions of the vertices in the beampipe, and information about the quality of each fit is saved into many variables.

Particles may also have associated ECAL or HCAL energy deposits, which are reconstructed as calorimeter clusters. For neutral particles which leave no hit in the tracker, this the only way of identifying them, but calorimeter clusters also improve the energy reconstruction for charged particles with poorly reconstructed track parameters, such as electrons and their high

1. The beam spot is the region where the two proton beams counter-rotating in the LHC collide inside the CMS detector.

bremsstrahlung rates, and low quality or high p_T charged hadron tracks. First, so called topological clusters are constructed from seeds which have an energy higher than a given threshold and higher than the neighboring calorimeter cells. Neighboring cells are added to the topological cluster as long as their energy is higher than twice the noise level, with more stringent requirements in the endcaps where noise levels are higher [140]. This topological cluster may contain many seeds and thus the energy deposits of many different particles. In a topological cluster with N seeds, an iterative optimization algorithm based on a Gaussian mixture model assigns a best-fit value to the energies and positions of the N seeds, which are then considered different clusters coming from different particles.

The PF algorithm reconstructs events in a specific order, with the easiest objects being reconstructed first: muons, then electrons and isolated photons, hadrons and non-isolated photons, and finally hadrons which undergo nuclear interactions in the tracker. These objects are reconstructed using the tracks and/or the calorimeter clusters mentioned in this section. A final post-processing step checks that events do not have artificially large missing E_T , which can be due to cosmic muons going through the detector, muons with a wrongly-reconstructed p_T , or misidentified particles. This reprocessing is not performed for non-PF objects.

5.2 Muons

The high mass analysis in this thesis uses an algorithm optimized for high p_T muons rather than the default PF algorithm, but the study of the tau identification efficiency presented in section 5.7 does make use of PF muons. Both kinds of muon selection are thus explained in this section.

There are three types of muon tracks, classified according to how they are reconstructed [144]:

- **Standalone muons** are built with information from all muon subdetectors (RPCs, DTs and CSCs) with a Kalman filter technique. The trajectory reconstruction is seeded by groups of track segments in either the DTs or the CSCs.
- **Tracker muons** are reconstructed inside out in the second-to-last step of the CMS iterative tracking algorithm. Each track passing minimum p_T (> 0.5 GeV) and total momentum (> 2.5 GeV) selections is extrapolated to the muon chambers. If the extrapolated trajectory matches a track segment in the muon system, then the inner silicon track is considered to be a tracker muon.
- **Global muons** are built outside-in in the last step of CMS iterative tracking. Standalone and inner tracks are extrapolated to a common surface, and are assigned to the same global muon if they match. Then a combined fit of the track parameters is performed, using information from both the muon system and the silicon tracker. If a global muon shares the same inner track as a tracker muon then the two candidates are merged.

Muon track reconstruction is extremely efficient, with over 99% of muons produced within the geometrical acceptance of the CMS detector being reconstructed as global or tracker muons.

Muons with low p_T may not have enough energy to go through all layers of the muon system and are thus more likely to be reconstructed as tracker muons. However, hadrons may also occasionally punch through the HCAL and trigger some activity in the first muon station, the tracker muon collection is thus contaminated by charged hadron fakes. On the other hand, standalone muons are contaminated by cosmic muons which are completely unrelated to the collision event.

The number of fake muons can be reduced by using certain identification criteria [144]:

- The χ^2 of the track fit.
- The number of hits per track, which can be divided into hits in the muon system, hits in the inner tracker and even specifically hits in the pixel detector.
- Compatibility with the event primary interaction vertex, which is the reconstructed vertex with the highest sum of p_T^2 from physics objects. (There can be many proton-proton interaction vertices, and it is extremely rare for more than one of them to produce a high p_T muon.)
- Compatibility between the inner track propagated outwards and track segments reconstructed in the muon system.
- A kink-finding algorithm searches for sudden changes in direction of the inner muon track, by splitting the track in two at various points and using a χ^2 score to evaluate the compatibility of the two parts of the track.
- For global muons, the compatibility of the inner track to the muon system track.

These criteria can be combined in different ways for different trade-offs in muon identification efficiency versus rejection of fake muons. Some criteria are optimized to identify muons from decays in flight, others only allow muons originating from the primary vertex. Most relevant to this thesis are the “medium” and “high p_T ” muon identification algorithms. The medium ID is used in the study of the tau identification efficiency, while the high p_T ID is optimized for muons with $p_T > 200$ GeV and is used in the high mass search for new physics.

A medium muon is a PF muon reconstructed as a tracker or a global muon, with an inner track with hits in more than 80% of the layers it goes through. A tighter or looser requirement on the compatibility of the muon system track segments with the inner track is required according to whether the muon is a tracker-only or global muon. This looser muon segment compatibility requirement for global muons is compensated by extra requirements on the goodness-of-fit of the global track, on the matching of the inner and standalone muon tracks and the χ^2 score calculated by the kink-finding algorithm [144].

The high p_T muon ID had some of its requirements relaxed compared to what is described in [127] in order to recover 2-3% efficiency at the highest momenta; the changes are explained in [145]. That said, a high p_T muon is reconstructed as a global muon, with at least one muon chamber hit, hits in at least six layers of the inner track, and at least one pixel hit, guaranteeing a minimum quality on the p_T measurement and suppressing muons coming from decays in flight. The quality of the p_T measurement is also ensured by requiring the relative error on the muon p_T track to be less than 30%. To further suppress muons from decays in flight, as well as cosmic muons and muons originating from other interaction vertices, the inner track must be compatible with the primary interaction vertex: the transverse impact parameter d_{xy} of the track with respect to the primary vertex needs to be below 2 mm, while the distance d_z between the two along the z -axis must be less than 5 mm. The final requirement of the high p_T ID is that muon track segments compatible with the inner track must be present in at least one station other than the first station. Segments in the first station could come from hadronic punchthrough from the HCAL, especially if no segment is found in the other stations. Tracks with compatible segments in the first station are fine as long as there are other compatible segments in other stations.

One variable is commonly used to distinguish prompt muons originating from the primary interaction vertex from non-prompt muons produced by hadrons decaying within a jet: the muon isolation, which characterizes how much activity there is in the detector in a close neighborhood around the muon. When there is a lot of activity around the muon, the isolation score

is high and the muon is likely part of a jet. This is why analyses interested in prompt muons, such as the tau efficiency study or the high mass LFV analysis, require muons to stay below a maximum isolation threshold. There are two different strategies for computing isolation: one uses only reconstructed tracks and thus does not take into account neutral particles (tracker-based isolation, used in the LFV analysis), and another uses PF charged hadrons and neutral particles (PF isolation, used in the tau efficiency study). When computing the PF isolation, charged hadrons are only taken into account if they originate from the primary interaction vertex. Estimating what fraction of the neutral particles originates from the primary vertex is not as straightforward because they leave no track. The isolation contribution from neutral particles originating from pileup vertices is estimated by multiplying the charged hadron pileup activity by a corrective factor, in a similar way to that described for the tau isolation in section 5.7.

The muon momentum is determined by the “Tune-P” algorithm, which selects one of four different momentum calculations according to the relative error on the muon p_T . There is an inner-tracker-only fit, good for low p_T muons, a fit with the inner tracker and the first muon station, a “picky fit” which deals with multiple hits in a single chamber (which likely come from a particle shower), and finally a “dynamic truncation” fit which deals with muons experiencing large energy losses and thus changes in their trajectories. For PF muons, the p_T assignment is further tuned in a post-processing step meant to cut down on artificially large missing E_T [144].

5.3 Electrons

Electrons are charged particles and undergo large electromagnetic interactions with matter. In CMS, they deposit almost all of their energy in the ECAL, which was made for this very purpose, and they also leave hits in the tracker, distinguishing them from photons. Electron reconstruction is not straightforward due to the already mentioned high interaction rates with matter: they are likely to start a shower by emitting bremsstrahlung photons² as they go through the tracker material, kinking their trajectories and spreading out in the ϕ direction their energy deposits in the ECAL. Closely spaced ECAL energy clusters can thus be grouped together into superclusters corresponding to what was originally a single electron (or photon) which triggered a shower while still traversing the tracker [140]. Electron tracks are reconstructed by a dedicated Gaussian-Sum Filtering (GSF) algorithm.

Superclusters are designed to take into account the bending of electrons in the CMS magnetic field, by being wider in ϕ than in η . The $\Delta\eta$ and $\Delta\phi$ depend on E_T because higher energy electrons have straighter trajectories. There is a first “mustache” algorithm which constructs superclusters using only information from the ECAL and preshower, and which is used to seed the reconstruction of electrons and photons. The mustache superclusters seed a second, “refined” algorithm, which makes use of tracker information to extrapolate the paths of bremsstrahlung photons and the tracks of converted electron pairs, and then decides if a given cluster should belong to a supercluster or not. This refined algorithm is what ultimately determines any ECAL-related quantities of electrons and photons.

Unlike the Kalman filter, which uses a single Gaussian as its probability function, the Gaussian-Sum Filtering (GSF) algorithm allows a mixture of multiple Gaussian functions, which provides a better approximation of the energy loss function [146]. However, the use of this more complex function makes the track extrapolation more computationally expensive, which means

2. Which may themselves convert to electron pairs.

it cannot be run over all hits in the tracker. The track seeding has to be constructed carefully in order to keep the efficiency high while rejecting enough hits for the computing expense to be minimized. There are two types of seeding, one based on ECAL “mustache” superclusters and the other on tracker hits. The ECAL seeding requires a supercluster with $E_T > 4 \text{ GeV}$ and with no significant energy deposits behind it in the HCAL: in a $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2} < 0.15$ cone centered on the supercluster, the HCAL energy needs to be less than 15% of the supercluster energy. The energy-weighted average position of the clusters which together form the supercluster should be located on the helical trajectory that the original electron would follow if it did not interact with the tracker; this fact is used to propagate backwards the electron candidate’s trajectory for both positive and negative charge hypotheses, and then check if the extrapolation matches any collection of hits in the innermost layers and disks of the tracker. The ECAL-driven seeds are any matched doublets or triplets of hits in the barrel pixel layers, the endcap pixel layers or the tracker endcaps. The tracker-based seeding is the same process in reverse. It uses tracks with $p_T > 2 \text{ GeV}$ reconstructed by the already-described CMS iterative tracking algorithm. These tracks are not specific to electrons, a multivariate analysis (MVA) checks if any of these general purpose tracks matches any supercluster. The combined efficiency of the two types of seeding is over 95% for electrons originating from a Z boson decay [147].

The GSF tracking algorithm starts from the seeds and proceeds iteratively through each tracker layer, resembling a set of multiple Kalman filters working in parallel [146]. The algorithm extrapolates the trajectory to the next layer, with relatively loose criteria, making it possible to follow the electron trajectory even after a bremsstrahlung emission. Multiple hits can be found to be compatible with the trajectory in a given layer, in that case multiple candidate trajectories are created and proceed to the next iteration step, with a maximum of five candidates per layer. Up to one missing hit is allowed for a valid trajectory candidate, but a large χ^2 penalty is applied on candidates missing a hit in order to suppress hits coming from bremsstrahlung photons which then convert to electron pairs. Once the final set of track candidates is completed, a GSF fit of the track parameters is performed, using a mix of Gaussian distributions to approximate the energy loss in each layer [147].

Refined superclusters are constructed to recover any missed bremsstrahlung photon, including those that convert to electron pairs later. The algorithm first matches each GSF track to any PF cluster directly in its trajectory at the exit of the tracker. It also extrapolates a straight line tangential to the GSF track for each tracker layer, which would potentially correspond to a bremsstrahlung photon, and checks if that straight line matches any PF cluster. To identify clusters corresponding to bremsstrahlung photons converted to electron pairs, an MVA algorithm based on displaced Kalman filter tracks is used.

A variety of strategies are used to discriminate prompt isolated electrons³ from backgrounds such as photon conversions, misidentified jets and electrons originating from hadron decays. Discriminating variables include observables characterizing how well the superclusters and tracks match, calorimetric observables such as the transverse shape of the electromagnetic shower and the energy fraction deposited in the HCAL, and tracking observables exploiting the difference between GSF and Kalman filter tracks. For relatively low energy electrons, both a simple strategy with sequential selections on discriminating variables and a machine-learning-based MVA strategy are used [147]. To identify high energy electrons, such as those in the final state of the high mass LFV analysis, a dedicated set of sequential selections is applied, called HEEP ID, for high energy electron pairs identification. The HEEP ID is not compatible with the PF framework.

3. These would be electrons originating from the primary interaction vertex (i.e. prompt) and which are not surrounded by a lot of nearby activity in the tracker or calorimeters (i.e. isolated).

The HEEP ID selects electrons differently according to whether the corresponding supercluster is in the barrel or one of the endcaps, to account for the different geometries in the two regions of the detector. A “crack” region between the two, which has poor ECAL coverage, is vetoed. The minimum transverse energy E_T is 35 GeV, and the GSF track seed needs to be ECAL-driven. Many other selections are used, summarized in table 5.1 and in the list below [148]:

- The track η extrapolated from the measures in the inner tracker layers must match the η of the supercluster seed. The difference between these two pseudo-rapidities is called $\Delta\eta_{\text{in}}^{\text{seed}}$, and the selection on it is tighter in the barrel (< 0.004) than in the endcaps (< 0.006) because there is more material in the way of an endcap electron, and the precision on the η measurement is reduced.
- $\Delta\phi_{\text{in}}$ is a similar variable to $\Delta\eta_{\text{in}}^{\text{seed}}$, but it corresponds to a difference in azimuthal angles rather than in pseudo-rapidity. The selection on the ϕ difference is significantly looser ($\Delta\phi_{\text{in}} < 0.06$) than for η because of bremsstrahlung photon emissions by the electron causing the ϕ energy distribution to be broader.
- In order to reduce the number of hadron fakes, there is a selection on the relative amounts of energy deposited in the HCAL and ECAL, H/E , which depends on the ECAL energy. The HCAL energy is integrated in a cone of radius 0.15 around the electron’s position in the calorimeter, while the ECAL energy is simply the energy of the electron supercluster.
- The η energy spread of the endcap ECAL showers must not be too large. The spread is measured by the $\sigma_{i\eta i\eta}$ variable, calculated on a 5×5 crystal block centered on the seed crystal. The formula for $\sigma_{i\eta i\eta}$ is:

$$\sigma_{i\eta i\eta} = \sqrt{\frac{\sum_{i \in 5 \times 5} (\eta_i - \bar{\eta})^2 w_i}{\sum_{i \in 5 \times 5} w_i}}, \quad w_i = \max\left(0, 4.7 + \log\left(\frac{E_i}{E_{5 \times 5}}\right)\right) \quad (5.1)$$

where i is an index running over the ECAL crystals in the 5×5 block, $\bar{\eta}$ is the supercluster η and $E_{5 \times 5}$ is the energy of the crystal block.

- In the barrel, the η energy spread within the 5×5 crystal block is calculated differently, by taking the ratio of the energy deposited in a narrow η strip around the seed crystal to the total energy deposited in the 5×5 block. Two such ratios are calculated, $\frac{E_{1 \times 5}}{E_{5 \times 5}}$ and $\frac{E_{2 \times 5}}{E_{5 \times 5}}$. $E_{1 \times 5}$ is the energy deposited in a 1×5 strip in the $\eta \times \phi$ plane, centered on the seed crystal. For $E_{2 \times 5}$, the choice of strip is ambiguous since there are two 2×5 strips centered on the seed crystal that can be chosen. The ambiguity is resolved by choosing the most energetic strip. The HEEP selection in the barrel is a logical OR of selections on $\frac{E_{1 \times 5}}{E_{5 \times 5}}$ and $\frac{E_{2 \times 5}}{E_{5 \times 5}}$.
- The number of lost hits in the innermost layers of the tracker, before the first GSF hit, needs to be at most 1. This suppresses converted photons and ensures track quality.
- There is a requirement on the transverse impact parameter d_{xy} to select prompt rather than displaced electrons. d_{xy} is the distance of closest approach of the GSF track to the primary interaction vertex in the transverse plane. Due to the poorer track reconstruction in the endcaps, the d_{xy} selection is looser there than in the barrel.
- The HEEP electron needs to be a standalone particle, with little nearby activity in the detector. A calorimetric isolation variable is calculated for the electron, summing the isolations in the ECAL and the first layer of the HCAL (HCAL Depth 1). The ECAL isolation is the scalar E_T sum of all ECAL crystals (above a minimum threshold: 80 MeV in the barrel and 100 MeV in the endcaps) in a cone of $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2} = 0.3$ around the electron, excluding those in an inner cone of radius 3 crystals and those in a

three η layers wide strip centered on the electron. The HCAL Depth 1 isolation is the scalar E_T sum of all towers in the first layer of the HCAL, in a cone of $\Delta R = 0.3$ around the electron, excluding an inner cone of $\Delta R = 0.15$. The point of excluding an inner cone or an inner strip from the isolation is to remove energy deposits corresponding to the electron. The calorimetric isolation, i.e. the sum of the ECAL and HCAL Depth 1 isolations, is strongly dependent on the electron energy due to the way electromagnetic showers spread at higher energies, and this is why the selection on the calorimetric isolation has a linear dependence on the electron E_T . The selection also depends on ρ , which is the median of the E_T density in the event per unit area, and includes energy deposits from pileup interactions as well as those from the primary interaction. As the pileup increases, so does ρ . Particles coming from pileup interactions can contribute to the calorimetric isolation, so relaxing the isolation requirement as ρ increases is a way to keep the selection efficiency high as pileup increases.

- Another isolation variable is calculated, the track isolation, which is defined as the scalar p_T sum of all reconstructed tracks in between an inner ($\Delta R = 0.04$) and outer ($\Delta R = 0.3$) cones around the GSF electron. The inner cone is excluded in order to remove tracks corresponding to the electron itself. Only tracks above a minimal p_T threshold (700 MeV) and with $\Delta z < 0.2$ cm with respect to the GSF track are considered. The Δz selection removes tracks coming from pileup interactions.

Variable	Barrel	Endcap
η range	$ \eta_{SC} < 1.4442$	$1.566 < \eta_{SC} < 2.5$
E_T	$> 35 \text{ GeV}$	$> 35 \text{ GeV}$
isEcalDriven	$= 1$	$= 1$
$ \Delta\eta_{\text{in}}^{\text{seed}} $	< 0.004	< 0.006
$ \Delta\phi_{\text{in}} $	< 0.06	< 0.06
H/E	$< 1/E + 0.05$	$< 5/E + 0.05$
full $5 \times 5 \sigma_{\text{in}\eta}$	-	< 0.03
full $5 \times 5 E^{2 \times 5} / E^{5 \times 5}$	$> 0.94 \text{ OR } E^{1 \times 5} / E^{5 \times 5} > 0.83$	-
ECAL + HCAL Depth 1 isolation	$< 2 + 0.03 \cdot E_T + 0.28 \cdot \rho$	$< 2.5 + 0.28 \cdot \rho \text{ for } E_T < 50 \text{ GeV} \text{ else } < 2.5 + 0.03 \cdot (E_T - 50) + 0.28 \cdot \rho$
Track Isolation	< 5	< 5
Inner Layer Lost Hits	≤ 1	≤ 1
$ d_{xy} $	< 0.02	< 0.05

Table 5.1 – Electron HEEP ID (v7.0) selections, used for 2016-17 data. Note that ρ is the median of the E_T density in the event per unit area, it includes energy deposits from pileup interactions as well as those from the primary interaction. In 2018, a slightly modified ID is used (V7.0-2018Prompt).

The selections just described correspond to HEEP ID v7.0, used in 2016-17 data. A revised ID V7.0-2018Prompt is used in 2018, with two changes in the endcap selections. This tuning was necessary because in 2018 some modules of the HCAL Endcap Minus (HEM) were non-operational during part of the data-taking (from Run 2018C onwards). The problematic region is limited in η and corresponds to the HEM 15-16 region of the HCAL. Due to this issue the HEEP identification scale factors (i.e. the ratios between the HEEP efficiency in data and in simulations) were no longer constant as a function of the p_T of the electron. It is important for the scale factors to be constant for data analyses using high energy objects because there are few data events at high energy and it is thus hard to be confident about extrapolating a trend observed at low energy to much higher energies. That said, the following changes to the HEEP ID produced constant scale factors in 2018:

- The endcap H/E selection is changed to $H/E < (-0.4 + 0.4|\eta|) \cdot \rho/E + 0.05$.
- The endcap maximum ECAL + HCAL Depth 1 isolation is set to $2.5 + (0.15 + 0.07 \cdot |\eta|) \cdot \rho$ for $E_T < 50$ GeV, else $2.5 + 0.03 \cdot (E_T - 50) + (0.15 + 0.07 \cdot |\eta|) \cdot \rho$.

5.4 Photons and hadrons

Photons and hadrons are relevant to tau lepton identification since they are decay products of the tau. Taus may decay semi-leptonically into a tau neutrino and one or many hadrons. Tau decays do not directly include photons, but they do include neutral pions, which decay with almost 100% branching ratio to photon pairs.

Photons are similar to electrons in that they deposit most of their energy in the ECAL, and they can also shower in the tracker, by converting to electron-positron pairs. The ECAL superclusters explained in the previous section are also used for photon reconstruction. Isolated photons, with no nearby activity in the tracker or calorimeters, are identified first: they are likely not electrons as they do not have any associated GSF tracks, and a selection on the HCAL energy fraction eliminates most of the hadronic background. Non-isolated photons are more difficult to distinguish from hadrons, which is why these two types of particles are identified at the same time.

ECAL and HCAL clusters not associated to any tracks are assigned to photons and neutral hadrons. Within the tracker acceptance ($|\eta| < 2.5$ in 2016 or 2.7 in 2017-18), which is the region of the detector relevant to this thesis, all of these ECAL clusters are assigned to photons and the HCAL clusters to neutral hadrons. This can be safely done because neutral hadrons deposit only 3% of their energy in the ECAL [140].

The clusters which are linked to tracks are treated differently (note that HCAL and ECAL clusters may be linked to each other as well as being linked to a track). The cluster energies are first recalibrated and compared to the corresponding track momenta. Clusters with linked tracks are likely to contain charged hadrons, but there could also be other objects. If the total calorimetric energy is significantly higher than the energy derived from the tracker, the excess is assigned first to photons if there is an ECAL cluster, and then any remaining energy is assigned to neutral hadrons. If the calorimetric energy matches the track momenta, then no photon or neutral hadron is created, and instead the charged hadron energies and momenta are recalculated using both tracker and calorimeter information. More rarely, the calorimetric energy can be lower than that implied by the tracker, and in that case a search for global muons with relaxed criteria is performed. If that fails to get rid of the discrepancy, a search for misreconstructed tracks with relatively large p_T uncertainties is done, and any such track found is removed from the PF collection. The procedure is repeated until no such track remains, or until the calorimetric energy is equal to or higher than the tracker-derived energy. These two cases are treated as already described in this paragraph [140].

5.5 Jets

Jets are clusters of closely packed particles. Particles are often clustered in this way because a jet typically originates from a single quark or gluon (quarks and gluons can be collectively referred to as partons). Interactions between the constituent partons of two colliding LHC protons may lead to the creation of new quarks or gluons, which carry color charges and are never directly detected. Instead, they lead to the creation of one or more hadrons, which may be

unstable and decay to other particles. The result is that a single parton produced during the collision may lead to a collection of many different final state particles. The point of reconstructing jets is to have some information about the original partons.

PF jets are built with the anti- k_T algorithm. The anti- k_T jet clustering algorithm was developed to be free of both infra-red (low-energy) and collinear divergences, and with the property that the jet shapes are not affected by the emission of low-energy particles (this property is called soft resilience) [149]. The distribution and the shapes of the jets are however influenced by high-energy particles. The algorithm is iterative, in each step it calculates “distances” d_{ij} between entities i and j , and distances d_{iB} between the entity i and the beam. Entities can be either particles or collections of particles grouped by a previous iteration of the algorithm. If the smallest distance is d_{ij} then i and j are merged into a single entity and the distances are recalculated; if the smallest distance is d_{iB} then i is declared a jet and removed from the list of entities. The algorithm stops when there are no more entities left. The distances are defined as follows:

$$d_{ij} = \min \left(\frac{1}{k_{Ti}^2}, \frac{1}{k_{Tj}^2} \right) \frac{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}{R^2},$$

$$d_{iB} = \frac{1}{k_{Ti}^2} \tag{5.2}$$

where k_{Ti} , y_i and ϕ_i are respectively the transverse momentum, rapidity and azimuthal angle of particle i . R is a free parameter determining the size of the jets, the higher it is, the bigger the jets. In CMS, its value is $R = 0.4$ for AK4 jets and $R = 0.8$ for AK8 jets. Note that only the transverse momentum of the highest momentum entity contributes to d_{ij} , ensuring soft resilience.

Jets can be contaminated by particles originating from pileup interactions. To mitigate their contributions, the charged hadron subtraction (CHS) technique is used. Charged particles whose tracks are associated to secondary vertices rather than the primary interaction vertex are removed from the list of particles available to reconstruct jets and other objects [150]. Neutral particles do not leave tracks, their pileup contributions are mitigated with an event-by-event, jet-by-jet method based on jet areas [151].

The calibration of jet energies and the measurement of the jet energy resolution are performed by techniques described in [152]. CHS PF jets are used as seeds to reconstruct tau leptons which decay into hadrons and a neutrino.

Jets originating from b quarks are identified by the Run 2 versions of the combined secondary vertex (CSVv2 in 2016, DeepCSV in 2017-18) algorithm, which exploits characteristics of jets coming from heavy quark flavors (such as b quarks) as opposed to those coming from light flavors. B hadrons have a lifetime of about 1 ps, which means that there is a secondary vertex at the point where they decay, 1-10 mm away from the primary interaction vertex. The higher mass of b quarks also means that the resulting jets are less collimated, and the decay products of B hadrons have a relatively large track momentum perpendicular to the jet axis. Many variables describing secondary vertices and the track p_T and η relative to the jet axis are used to train a neural network to discriminate between b jets and light flavor jets [153]. Events containing b quark jets are rejected in the tau identification study presented in section 5.7, and they are used to construct a control region enriched in top-antitop events for the high mass LFV analysis.

5.6 Missing transverse energy

In CMS, protons about to collide have negligible momentum in the transverse plane, i.e. the plane that is orthogonal to the beam. Since momentum is conserved, the total transverse momentum is also negligible after a collision. The momentum is also conserved along the z -axis but this information cannot be used in practice because proton remnants are lost in the beam pipes after each collision. There is no such complication in the transverse plane: if the vector sum of the \vec{p}_T 's of all particles is not zero, then either a measuring error happened, or there were one or many invisible particles in the final state. A missing transverse momentum vector is defined to account for any discrepancies between the measured total \vec{p}_T of an event and its theoretical value of zero. The missing \vec{p}_T vector of a given event is then simply the negative \vec{p}_T sum of all particles in the event, and the magnitude of the vector is called missing transverse energy (missing E_T , \cancel{E}_T or MET⁴). For analyses targeting final states with neutrinos, such as the search for high mass LFV in the tau channels, the MET is important for estimating the energy and momentum of these invisible particles. The MET used in the analysis is the PF MET, i.e. using the negative \vec{p}_T sum of all particles constructed by the PF algorithm.

When jet energies are recalibrated [152], these corrections are propagated to the MET in order to make the event description consistent. There are however other sources of MET mismeasurement, and there are dedicated filters to remove events with spuriously high missing energy caused by detector malfunction or failures in reconstruction [154]. The recommended MET filters for high mass analyses are:

- **HCAL noise filters**, which remove events with anomalously high MET due to particles hitting the electronics or with channels firing randomly without any particle going through them. A dedicated filter combines HCAL, ECAL and tracker information to remove events with anomalous isolated HCAL activity.
- **ECAL noise filters**. Most of the ECAL noise can be corrected during event reconstruction, but some events cannot be corrected and are best removed by dedicated filters. One filter removes events from specific noisy ECAL supercrystals. A second filter considers the behavior of some dead channels in the ECAL. The dead channels do not provide energy measurements for their corresponding ECAL tower, however for most of them the trigger primitives are available and they allow the energy deposited in the tower to be recovered. The trigger primitives record a narrower range of energies than the ECAL readout system, so the energy measured by the trigger primitives can saturate. The filter removes events in which the energy in one of these dead cells approaches the saturation level.
- **A beam halo filter**, which removes events in which halo particles interact with the detector. Halo particles fly parallel to the beam, and are produced when e.g. a proton interacts with a particle in the beampipe outside of the collision region (no vacuum is perfect and there are still a few other particles in the beampipe). High energy halo muons can interact with the calorimeters and produce entirely spurious MET. The filter uses information from the CSCs and the calorimeters to identify and remove such events.
- **Reconstruction filters**, which remove events with anomalously high MET due to misreconstructed PF muons.

The resolution and the scale⁵ of MET measurements are estimated by considering events with a high p_T Z boson or photon recoiling against one or many jets. When e.g. the Z boson decays to a dimuon pair, there is no true MET in the event, but it is possible to simulate MET

4. In this thesis, the MET shorthand also refers to the missing transverse momentum vector, as conventionally done in the CMS collaboration.

5. The scale determines if there is any systematic bias between the measured and the “true” MET.

by pretending that the dimuon pair is not part of the event. The event’s “missing \vec{p}_T ”, calculated using just the hadronic activity, should correspond to the actual dimuon \vec{p}_T measured in the same event. By comparing the distribution of directly measured Z boson transverse momenta to the distribution of those transverse momenta reconstructed from hadronic activity, it is then possible to estimate the MET scale and resolution. For boson $p_T > 100$ GeV, the scale is close to 1, with small deviations due to imperfect calibration of the hadronic energy scale. For $p_T > 200$ GeV, the resolution on the two components of the \vec{p}_T^{miss} vector is $\sim 9\text{-}13\%$ [154].

Using the PF MET for the high mass LFV analysis is inconsistent because, as explained in sections 5.2 and 5.3, the electrons and muons used in the analysis are not PF. The MET needs to be recalculated in order to keep a consistent event description. For an event with a non-PF high- p_T muon, this is done by first subtracting the \vec{p}_T of the high- p_T muon from the PF MET, and then checking if the high- p_T muon is compatible ($\Delta R < 0.2$) with any PF muon in the same event. If it is, then the closest matching PF muon \vec{p}_T is added to the MET⁶. The equation relating the corrected missing transverse momentum vector $\vec{p}_T^{\text{corrected}}$ to the PF missing transverse momentum \vec{p}_T^{PF} is then:

$$\vec{p}_T^{\text{corrected}} = \vec{p}_T^{\text{PF}} - \vec{p}_T^{\text{high } p_T \mu} + \vec{p}_T^{\text{PF } \mu, \text{ matched}} \quad (5.3)$$

This corrected PF MET is the one used in the LFV analysis. Most of the time the correction is small but it can be large in events where a high p_T muon is identified but not a PF muon (the high p_T muon selections are looser). The same correction procedure is done with HEEP electrons in the $e\tau$ final state, while nothing needs to be done in the $e\mu$ channel because MET is not a factor there.

5.7 Taus

Taus are unstable particles, with a 1.777 GeV mass and a mean lifetime of 2.8×10^{-13} s [155]. Unlike muons, this lifetime is too short for them to cross the detector without decaying: even a 200 GeV tau, with a relativistic γ factor of more than 100, would typically travel less than 1 cm before decaying. As a comparison, the LHC beampipe has a diameter of 5.5 cm. As shown in table 5.2, about one third of tau decays are fully leptonic with final states including two neutrinos and either a muon or an electron. No specific effort is made at the particle identification level to check if a given electron or muon originates from a tau decay. About two thirds of the time, taus decay semi-leptonically into a tau neutrino and one or many hadrons. These taus are called hadronic taus (τ_h), and there are dedicated algorithms to reconstruct and identify them.

5.7.1 Reconstruction of hadronic taus

The reconstruction of τ_h candidates is seeded by charged hadron subtracted PF jets with $R = 0.4$, and depends on the decay mode, i.e. how many charged hadrons and neutral pions there are in the final state. No particular treatment is needed for charged hadrons, they can be taken directly from the PF reconstruction described in section 5.4. However, neutral pions are unstable, they decay almost immediately and with close to 100% probability to two photons, which are themselves likely to convert to electron pairs while traveling through the tracker material. The electron and positron of each pair have their trajectories bent in opposite directions

6. Remember that the MET is a negative \vec{p}_T sum, removing a particle from the equation means actually adding its \vec{p}_T back in.

τ^- decay mode	Resonance (mass in MeV)	Branching ratio (%)
$e^- \bar{\nu}_e \nu_\tau$		17.8
$\mu^- \bar{\nu}_\mu \nu_\tau$		17.4
$h^- \nu_\tau$		11.5
$h^- \pi^0 \nu_\tau$	$\rho^-(770)$	25.9
$h^- \pi^0 \pi^0 \nu_\tau$	$a_1^-(1260)$	9.5
$h^- h^+ h^- \nu_\tau$	$a_1^-(1260)$	9.8
Other decay modes with hadrons		8.0
All hadronic decay modes		64.8

Table 5.2 – Decay modes of the negatively charged tau lepton, with their respective branching ratios [155]. Hadronic modes include any decays with pions or kaons in the final state, though the overwhelming majority of hadrons are specifically pions. Charged hadrons are represented by the symbols $h^{+/-}$. The middle column shows resonances which may be intermediate steps in some of the hadronic decay modes [156]. For example, a tau can decay in the following way: $\tau^- \rightarrow \nu_\tau \rho^-(770) \rightarrow \nu_\tau h^- \pi^0$.

by the solenoid magnetic field, resulting in a π^0 signature which is more spread out in ϕ than in η . This is taken into account by the algorithm by reconstructing neutral pions as clusters of electrons and photons. These clusters tend to be wider in ϕ than in η and are hence called “strips”.

Strips are constructed iteratively, by first centering on the most energetic electron or photon in the jet. Then a search for other electrons or photons compatible with this first particle is performed: the second particle must have a ϕ - η position falling within a certain window around the original particle, and the size of the window is a decreasing function of both the first and the second particle’s p_T . A decreasing function is used because high p_T taus have more collimated decay products, and the corresponding strip size can be smaller than for lower p_T taus. When a strip contains more than one electron/photon, the strip p_T is defined as the sum of the p_T ’s of all strip constituents, and the strip’s ϕ - η position is calculated as the p_T -weighted average of the constituents:

$$\begin{aligned}
p_T^{\text{strip}} &= \sum_{i \in e, \gamma} p_T^i \\
\eta^{\text{strip}} &= \frac{1}{p_T^{\text{strip}}} \sum_{i \in e, \gamma} p_T^i \eta^i \\
\phi^{\text{strip}} &= \frac{1}{p_T^{\text{strip}}} \sum_{i \in e, \gamma} p_T^i \phi^i
\end{aligned} \tag{5.4}$$

In the next iterations, the compatibility window is calculated around ϕ^{strip} - η^{strip} and its size depends on p_T^{strip} as well as on the p_T of the potential new electron/photon addition to the strip. The algorithm stops when no further electron or photon can be added to the strip.

Reconstructed charged hadrons and strips can be combined into one of many different topologies of hadronic tau decays. These topologies are also known as decay modes, the four main decay modes are:

- One charged hadron (“1 prong + 0 π^0 ”).
- One charged hadron and one strip (“1 prong + 1 π^0 ”).

- Three charged hadrons (“3 prong + 0 π^0 ”).
- Three charged hadrons and one strip (“3 prong + 1 π^0 ”).

Neutral pions can fail to be reconstructed, so the above decay modes do not always correspond to what one would expect from table 5.2. The τ_h candidate is required to have a charge of ± 1 , its charge is computed by summing the charges of all its constituent charged hadrons. As a final step, a selection on the reconstructed tau mass is performed to guarantee some compatibility with the expected mass for each of the decay modes. The selection depends on the tau p_T and the change in the tau mass when the strips are included in the calculation [157].

5.7.2 Identification of hadronic taus

After being reconstructed, tau candidates still need to pass anti-muon, anti-electron and anti-jet discriminators in order to be properly identified as taus. This identification step has considerably evolved over the course of Run 2 in CMS. In the beginning of Run 2, all three discriminators (anti-muon, anti-electron and anti-jet) were different from the DeepTau algorithms that became available later and were ultimately used in the LFV analysis. Most of the progress of the high mass LFV analysis happened using the old tau identification algorithms, the more sophisticated DeepTau was a late addition to the analysis because it only became fully supported by the Tau Particle Object Group (Tau POG) at the end of 2019. Until the beginning of the same year, I was personally involved in the Tau POG, and my task was to determine the efficiency of the anti-jet discriminator by various versions of the old algorithms.

Before the DeepTau algorithm, an anti-electron discriminator making use of a boosted decision tree (BDT) was the standard algorithm used. The input variables of the BDT describe the distribution of ECAL energy deposits, the amount of bremsstrahlung emitted along the leading track, and the total number of particles, allowing for a distinction between electromagnetic and hadronic showers. Hadronic showers are of course those more likely to be associated with hadronic taus than with electrons. Multiple working points of the BDT anti-electron discriminant are available.

The old anti-muon discriminator was “cut-based”, rejecting τ_h candidates when there is little energy deposited in the ECAL and HCAL with respect to the momentum of the τ_h leading track, or when track segments are found in muon stations in a cone centered on the τ_h . A tight working point of the algorithm has the additional requirement that the same cone contain no hit in the outermost muon stations.

At the beginning of Run 2, two anti-jet discriminators were available: a simple “cut-based” one and a BDT-based algorithm. The cut-based anti-jet identification algorithm computes the isolation of the τ_h candidate by summing the p_T ’s of photons and charged particles which are close to the τ_h , excluding the particles which were used to construct it. Charged particles and photons could potentially originate from another interaction vertex, which did not produce the τ_h . For charged particles this pileup contribution is suppressed by requiring their tracks to originate from a region close to the interaction vertex ($d_z < 0.2$ cm), while the photon pileup contribution is estimated by summing the p_T ’s of charged hadrons incompatible with the τ_h production vertex ($d_z > 0.2$ cm), and multiplying that number by a correction factor ($\Delta\beta$) which takes into account the ratio of neutral to charged hadrons produced in proton-proton collisions [157]. The $\Delta\beta$ factor is chosen such that the τ_h identification efficiency is independent of the number of pileup vertices. The τ_h isolation I_τ can then be computed in the following way:

$$I_\tau = \sum p_T^{\text{charged}}(d_z < 0.2 \text{ cm}) + \max\left(0, \sum p_T^\gamma - \Delta\beta \sum p_T^{\text{charged}}(d_z > 0.2 \text{ cm})\right) \quad (5.5)$$

Higher values of the isolation correspond to more activity around the τ_h candidate, indicating that it is probably a fake. The cut-based algorithm simply accepts candidates with an isolation value lower than a given threshold, while rejecting the rest. There are several variations of the algorithm: different sizes of the region around the τ_h candidate where the isolation is computed, and different isolation thresholds with different trade-offs between the identification efficiency and the misidentification probability. The cut-based algorithm makes a final selection requiring that strips do not contain outlying electrons/photons which make up a significant fraction of the total p_T of the τ_h candidate.

The BDT anti-jet algorithm is based on a collection of decision trees, such as the one illustrated in figure 5.2 (left). A decision tree decides if a given τ_h is part of the signal or the background by starting from a root node and following branches according to whether certain variables are above or below given thresholds. The algorithm stops when it reaches the “leaf” at the end of a branch, which indicates whether the candidate is more likely to be signal or background. The outputs of many trees are used for the final decision on whether to accept or reject the τ_h candidate. The BDT is trained using simulated data of both signal and background processes. Signal samples include $Z' \rightarrow \tau\tau$ and $W' \rightarrow \tau\nu_\tau$ events for high p_T taus to also be well described. Events used for the training are reweighted so that the p_T and η distributions of the signal and background are identical, making the BDT decision independent of event kinematics. The discriminating variables used by the BDT include [157, 158]:

- The neutral and charged isolation sums of equation 5.5, as separate inputs.
- The reconstructed decay mode of the τ_h candidate.
- The transverse impact parameter d_0 of the τ_h candidate’s leading track, and the significance of d_0 , i.e. how (in)compatible with zero it is.
- A boolean indicating whether a decay vertex for the τ_h candidate has been reconstructed. If a secondary vertex is available, its distance with respect to the primary vertex is used as an input variable, as well as the significance of that distance.
- Variables describing the shape of the τ_h candidate, especially how spread out with respect to the τ_h axis the electron and photon constituents of the strips are.
- The total number of electrons and photons in a region close to the τ_h candidate.

Multiple iterations of this BDT algorithm were developed, with different training datasets. The first version had a misidentification rate equal to half of that of the cut-based algorithm for a similar identification efficiency [157].

The DeepTau algorithm is based on a deep neural network, and allows for an effective rejection of $\text{jet} \rightarrow \tau$, $\mu \rightarrow \tau$ and $e \rightarrow \tau$ fakes. As illustrated in figure 5.2 (right), the architecture of a neural network is more complex than that of a decision tree, with many possible cross-connections in the internal layers of the network. DeepTau takes as inputs many more variables than its predecessors. The inputs can be separated into high-level and low-level: the high-level variables include those which were used in the BDT algorithm as well as any other variable used for the τ_h candidate reconstruction, while the low-level variables include information from the inner tracker, the calorimeters and the muon system for all particle candidates reconstructed in a region close to the τ_h . For the same τ_h identification efficiency, the $\text{jet} \rightarrow \tau$, $\mu \rightarrow \tau$ and $e \rightarrow \tau$ misidentification rates are reduced by large amounts with respect to previous algorithms. The performance of the DeepTau algorithm can be compared to the old algorithms as follows [161]:

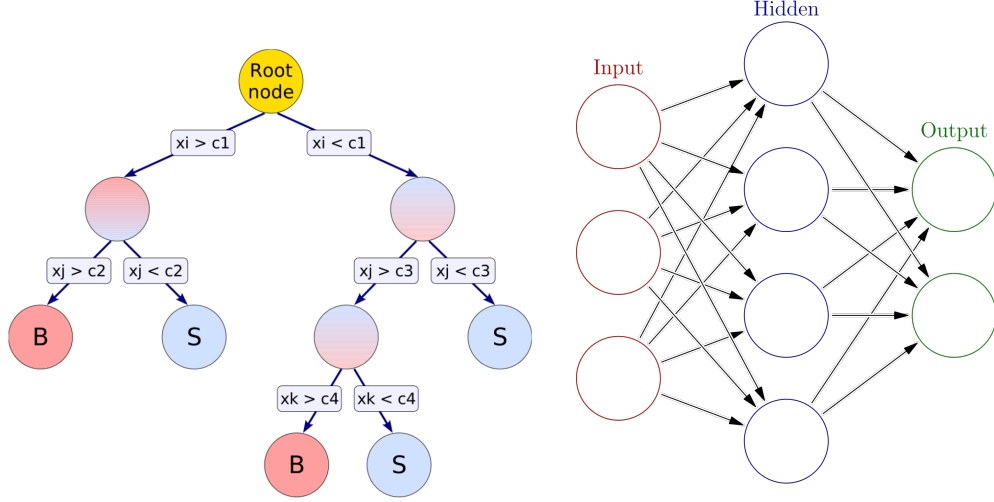


Figure 5.2 – Schematic illustration of different machine learning architectures. (Left) Example of a decision tree, with discriminating variables x_i , x_j and x_k and thresholds c_1 , c_2 , c_3 and c_4 [159]. At each node, the algorithm compares a variable to a given threshold, and moves along one branch or the other according to the result of the comparison. The procedure stops when a final node is reached, labeled either “B” for background or “S” for signal. Boosted decision trees (BDT) are made of a collection of such trees, with different weights for the outcomes of each tree. (Right) Example of a neural network: input variables are on the left and output variables on the right, with a hidden layer in between performing operations on the input layer variables to produce the outputs [160]. More complex architectures include several hidden layers, with each layer using as inputs the variables constructed by the previous layer.

- The DeepTau **anti-jet** discriminator efficiencies (misidentification rates) range from $\sim 88\%$ ($\sim 1 \times 10^{-2}$) for the loose working point (WP) to $\sim 75\%$ ($\sim 4 \times 10^{-3}$) for the tight WP. The efficiencies (misidentification rates) of the BDT-based algorithm are $\sim 85\%$ ($\sim 1.5 \times 10^{-2}$) for the loose WP and $\sim 67\%$ ($\sim 5 \times 10^{-3}$) for the tight WP⁷.
- The DeepTau **anti-electron** discriminator efficiencies (misidentification rates) are $\sim 94\%$ ($\sim 5 \times 10^{-3}$) for the loose WP and $\sim 78\%$ ($\sim 8 \times 10^{-4}$) for the tight WP, while for the BDT-based algorithm they are $\sim 88\%$ ($\sim 2 \times 10^{-2}$) for the loose WP and $\sim 79\%$ ($\sim 3 \times 10^{-3}$) for the tight WP.
- For the DeepTau **anti-muon** discriminator, the numbers are $\sim 99.8\%$ ($\sim 7 \times 10^{-4}$) for the loose WP and $\sim 99\%$ ($\sim 2 \times 10^{-4}$) for the tight WP, while for the old cut-based algorithm they are $\sim 99.2\%$ ($\sim 3 \times 10^{-3}$) for the loose WP and $\sim 98.7\%$ ($\sim 1 \times 10^{-3}$) for the tight WP.

The tau identification efficiencies are estimated using simulations of H bosons decaying to ditau pairs, the jet misidentification probabilities are derived from $t\bar{t}$ ⁸ simulations, and the electron and muon misidentification probabilities are determined using Drell-Yan simulations [161].

In this thesis, the only section that makes use of the old tau identification algorithms is the one which immediatly follows this paragraph, about the study of the anti-jet discriminator efficiency which I performed for the Tau POG. The LFV analysis uses the DeepTau discriminators. In part III of this thesis, the tau identification algorithms mentioned in the text are always DeepTau, except for one fake factor study presented in chapter 8.

7. Note that both even tighter and even looser working points exist for the DeepTau and the BDT algorithms.

8. This is the production of a top-antitop quark pair.

5.7.3 Study of the efficiency of the anti-jet discriminators

I worked for the Tau POG, and was tasked with determining the efficiencies of the various tau anti-jet discriminators in the data collected by CMS in 2017. More precisely, the point was to compute the ratio of the efficiency in data to the simulated efficiency, as well as any systematic error on that ratio. This ratio, or scale factor, can then be used by any CMS analysis using taus to reweight the simulations so that they better correspond to the expected data distributions, and the systematic error describes how precise that reweighting is. I stopped working for the Tau POG in early 2019, when I decided to focus my CMS service work on the trigger monitoring mentioned in chapter 4.

I did not work in the estimation of DeepTau identification efficiencies, but there were several versions of the cut-based and BDT-based algorithms to work on. There were three working points of the cut-based algorithm, and three different versions of the BDT algorithm, each with seven working points. There were also three different trainings of the BDTs: 2016v1, 2017v1 and 2017v2, with the last one improving on the other two. Scale factors for each of these discriminators had to be computed, as well as their dependence on the τ_h candidate's decay mode, p_T and η .

The scale factors were determined by using events from two different final states: $\mu\tau_h$ and $\mu\mu$. In the $\mu\tau_h$ final state, the goal is to maximize the relative number of Drell-Yan (DY) $Z \rightarrow \tau\tau$ events in which there is one hadronic tau. A $Z \rightarrow \tau\tau$ event can produce a $\mu\tau_h$ signature if one of the taus decays to hadrons and the other to a muon. Other tau decays could be considered, but the highest purity of $Z \rightarrow \tau\tau$ events is achieved with muons because CMS is particularly effective at identifying them, as indicated by what the “M” in CMS stands for. Purity in $Z \rightarrow \tau\tau \rightarrow \mu\tau_h$ events is further maximized by requiring strict selections on the muon, making it very unlikely that it is a fake. The identification of taus is not as effective at rejecting fakes as that of muons, so a considerable number of background events are still present in the $\mu\tau_h$ final state. However, in the $\mu\mu$ final state, almost 100% purity in $Z \rightarrow \mu\mu$ events can be achieved. The muon selections here are the same as in the $\mu\tau_h$ final state in order to make both data regions as similar as possible.

As mentioned in chapter 2, lepton universality dictates that the Z boson decays at the same rate to all charged lepton pairs. In particular, the branching ratios of $Z \rightarrow \mu\mu$ and $Z \rightarrow \tau\tau$ are identical. Thanks to lepton universality, in the case of a discrepancy between the data and the simulations in the $\mu\tau_h$ final state, the $\mu\mu$ final state can be used to constrain what part of that disagreement is due to a poor modeling of the production of the Z boson, which would also show up in $\mu\mu$ events.

The selections for the $\mu\tau_h$ final state are summarized in the list below:

- The trigger requires the presence of an isolated muon with $p_T > 27$ GeV.
- The muon is required to be identified by PF, pass the medium ID and have impact parameters compatible with the primary interaction vertex: $|d_{xy}| < 0.045$ cm and $|d_z| < 0.2$ cm. It must also have $p_T > 30$ GeV (to guarantee close to 100% trigger efficiency), $|\eta| < 2.4$ (that is the acceptance of the muon system), and a relative isolation < 0.15 .
- The tau needs to be a τ_h candidate for which a valid decay mode was found. It is required to have $p_T > 20$ GeV and $|\eta| < 2.3$, otherwise the decay mode finding algorithm does not work. A tight anti-muon and a very loose anti-electron vetos are required. Since there is a muon in this final state, rejecting muons is a higher priority than rejecting electrons, as many $Z \rightarrow \mu\mu$ events are expected. The final requirement is whatever anti-jet discriminator is being studied.

- The muon and tau are required to have opposite charges and to be well separated, with $\Delta R(\mu, \tau) > 0.5$. Events with one electron or with more than one muon are vetoed.
- There is a selection on the transverse mass m_T of the muon and the missing p_T vector: $m_T < 50$ GeV. The transverse mass of a collection of objects coincides with their invariant mass when all objects are massless and confined to the transverse plane (no longitudinal momentum). The invariant mass of a many-particles system is a useful quantity to calculate when we suspect that all the particles originate from the decay of the same heavy resonance. The invariant mass is defined in appendix A, it is a reconstruction of the heavy resonance mass. m_T is a reasonable approximation of the invariant mass when the invariant mass cannot be calculated, typically when one wishes to include the MET in the computation. m_T is defined as:

$$m_T = \sqrt{2 |\vec{p}_T^\mu| |\vec{p}_T| (1 - \cos \Delta\phi(\vec{p}_T^\mu, \vec{p}_T))} \quad (5.6)$$

where \vec{p}_T^μ and \vec{p}_T are respectively the muon p_T and MET vectors. m_T is explained in more detail in chapter 7, in the context of the high mass LFV analysis.

- A topological variable called P_ζ and describing how aligned the missing p_T vector is with the muon and the tau is used to discriminate signal from background. To calculate P_ζ , a unit vector $\vec{1}_\zeta$ is constructed halfway between the tau and muon p_T vectors, i.e. along the direction bisecting the $\Delta\phi$ angle between the two vectors. Then P_ζ is obtained by summing various momenta projected along this direction:

$$P_\zeta = \vec{p}_T \cdot \vec{1}_\zeta + 0.15 (\vec{p}_T^\tau + \vec{p}_T^\mu) \cdot \vec{1}_\zeta \quad (5.7)$$

where \vec{p}_T^τ is the p_T vector of the hadronic tau. When the MET is aligned in the same direction as the muon and the tau, as in signal $Z \rightarrow \tau\tau \rightarrow \mu\tau_h$ events, P_ζ tends to be positive. In this study, P_ζ is required to be higher than -25 GeV, keeping signal efficiency high while rejecting some background events where the MET direction is uncorrelated with the direction of the leptons.

- If multiple $\mu\tau_h$ pairs in the same event pass all requirements, then the one with the highest transverse momenta is selected.

For the $\mu\mu$ final state, the trigger and muon requirements are the same (except two muons are needed instead of just one). Otherwise, the muons are required to have opposite signs (OS), to be well-separated ($\Delta R(\mu, \mu) > 0.5$), and the dimuon invariant mass needs to be compatible with that of a Z boson resonance ($60 \text{ GeV} < m_{\mu\mu} < 120 \text{ GeV}$). Events with a third muon are vetoed.

Selections for both the $\mu\tau_h$ and $\mu\mu$ final states are summarized in table 5.3.

With these selections, the $\mu\mu$ final state is expected to contain almost only $Z \rightarrow \mu\mu$ events, while the $\mu\tau_h$ final state should contain the following non-negligible backgrounds:

- Drell-Yan (DY) $Z \rightarrow \mu\mu$ events with one $\mu \rightarrow \tau_h$ fake. These events are suppressed by the tight anti-muon discriminator applied on the τ_h .
- W+jets events, where one W boson is produced in association with one or more jets. One of the jets can be misidentified as a tau, while the W boson can decay into a muon and a neutrino. These events are suppressed by the requirement on the transverse mass of the muon and the missing p_T , which for W boson decays peaks around the W mass (80 GeV).
- QCD multijet events, in which several high p_T jets are produced by the proton-proton collision, one of which fakes a muon (or contains a non-isolated muon) and the other a τ_h . The strict muon selection requirements largely reduce the number of such events, but QCD cross sections are so high that a significant number of them still get through the selections.

	$\mu\mu$	$\mu\tau_h$
Trigger	IsoMu27	IsoMu27
Muon selection: PF muon, medium ID, $p_T > 30$ GeV, $ \eta < 2.4$, relative isolation < 0.15 , $ d_{xy} < 0.045$ cm, $ d_z < 0.2$ cm	Two OS muons with $60 < m_{\mu\mu} < 120$ GeV and $\Delta R(\mu, \mu) > 0.5$	One muon
Tau selection: $p_T > 20$ GeV, $ \eta < 2.3$, old DM finding algo, cut-based tight anti- μ discr., BDT very loose anti-e discr., anti-jet discr. (under study).	No tau	One tau, OS w.r.t. muon with $\Delta R(\tau, \mu) > 0.5$
$m_T(\mu, \cancel{p}_T) < 50$ GeV $P_\zeta > -25$ GeV	not applied not applied	applied applied
Extra lepton veto	Rejected events: events with more than two muons	Rejected events: events with (i) at least one electron or (ii) more than one muon

Table 5.3 – Selections used for the $\mu\mu$ and $\mu\tau_h$ final states. If multiple $\mu\tau_h$ pairs in the same event pass all requirements, then the one with the highest transverse momenta is selected.

- Less important backgrounds include events with top-antitop pairs ($t\bar{t}$) and diboson events (events with either WW, WZ, or ZZ). Tops decay with almost 100% probability to a b quark and a W boson, one of the W bosons of a $t\bar{t}$ event can decay into a genuine muon while the other decays into a genuine tau, or there could be one genuine muon while the tau is actually a misidentified jet (the proportion of fake muons is negligible). Similarly, diboson events can produce genuine muons and taus via W or Z decay, or produce one genuine muon while a jet is misidentified as a tau. The topological discriminator variable P_ζ suppresses $t\bar{t}$ events, but no particular measures are taken against diboson events.

The backgrounds with genuine muons and taus are well modeled by simulations. Some small corrections are necessary in order to resolve discrepancies with data for the efficiency of the trigger, the muon identification and the isolation requirements. These scale factors are computed by the Muon POG. The simulations also need to be reweighted so that their pileup distribution better matches that of the data. More details on the simulations and the pileup reweighting procedure are provided in chapter 6.

The tau energy scale in the simulations also needs to be corrected to better match the energy scale observed in data. These corrections are derived by the Tau POG. They start by selecting $e\tau_h$ and $\mu\tau_h$ final states, with $\mu\tau_h$ selections similar to those of table 5.3. The hadronic taus need to pass the very tight anti-jet BDT discriminant, in order to further increase sample purity. Templates are generated by using simulated events and varying the reconstructed tau energy between -6 % and +6 %, in steps of 0.1 %. A maximum likelihood fit is used to determine the template that better fits the data, and thus find the tau energy scale correction that needs to be applied to simulated events. Two different distributions are used for the fit: the τ_h mass, and the $\mu\tau_h$ or $e\tau_h$ visible mass. Energy scale corrections and their corresponding errors are obtained independently for each tau decay mode, and for genuine τ_h , $\mu \rightarrow \tau_h$ fakes and $e \rightarrow \tau_h$ fakes. They are all of order ~ 1 % [162].

Backgrounds with misidentified particles are not reliably modeled in the simulations. A data-driven approach is required in order to keep the uncertainties on the QCD multijet and W+jets distributions under control. At first, I estimated these backgrounds separately in a hybrid data-driven+simulations approach. I constructed a control region enriched in W+jets events by modifying the selection on the transverse mass m_T : instead of requiring $m_T < 50$ GeV, $m_T > 80$ GeV was required. In this region a scaling factor on the W+jets event yield could be calculated to make the number of events predicted by the simulations match the number of events observed in data. That same scaling factor could then be applied in the signal region ($m_T < 50$ GeV). This way, the normalization of the W+jets process was determined by data, while the shape of the mass distribution was still determined by simulation. The QCD multijet process shape and normalization were determined in a control region where the opposite sign (OS) selection on the $\mu\tau$ pair is flipped to a same sign (SS) requirement. This control region is almost entirely composed of equal parts QCD multijet and W+jets events, so by subtracting the (rescaled) W+jets mass distribution from the data mass distribution, the shape of the QCD distribution could be determined. The normalization was then corrected by an OS/SS scale factor⁹, calculated in yet another control region, in which the muon is anti-isolated and thus most likely a fake (this control region is made up almost entirely of QCD multijet events).

Figure 5.3 shows the $\mu\tau_h$ invariant mass distributions I obtained in both the W+jets region (before rescaling) and in the signal region, after estimating the W+jets and QCD multijets processes by the procedure described in the previous paragraph. The so-called visible mass is shown, which only uses the visible decay products of the tau and does not take into account the MET. The W+jets background is called “Electroweak” in the plot. The “Electroweak” histogram also includes diboson events, but these are less than 1 % of the total.

After producing the $\mu\tau_h$ signal region mass distribution plots, a fit of the signal and the various background processes to the data was performed, using the DY event normalization in the $\mu\mu$ final state as a constraint. The yield of the signal process ($\text{DY } Z \rightarrow \tau\tau \rightarrow \mu\tau_h$) was allowed to vary freely, while the backgrounds were only allowed to vary according to some nuisance parameters, defined by probability distribution functions. Large variations in background yield or shape with respect to expectations are disfavored when compared to smaller variations. The nuisance parameters are also applied on the signal distribution when they are relevant, but there is no constraint on the value of the ratio of the τ_h anti-jet discriminator efficiency in data to that same efficiency in simulations.

These fits behaved poorly when the W+jets and QCD multijet backgrounds were estimated separately by the already described hybrid data-driven/simulation procedure. Relatively large variations in the W+jets yield were compensated by equally large but opposite variations in the QCD multijet yield. This behavior hinted that lumping the QCD multijet and W+jet backgrounds together might produce better results. It was thought better to estimate together all processes involving misidentified τ_h in a completely data-driven manner by calculating so-called “fake factors” in much the same way as is done in the high mass LFV analysis. A region enriched in $\text{jet} \rightarrow \tau_h$ fakes is constructed by requiring the presence of τ_h candidates in the $\mu\mu$ final state, which are almost always jets. The fake factor is defined as the number of τ_h candidates passing the tau anti-jet discriminator divided by the number of candidates which do not pass the anti-jet requirement but do pass some baseline selections. The shape and normalization of the distributions of events with $\text{jet} \rightarrow \tau_h$ fakes is then obtained by applying the fake factor to events which pass all of the signal region $\mu\tau_h$ selections, except the τ_h candidate is required to fail the anti-jet requirement (but does need to pass the same baseline selections used to compute the fake factor). This procedure is explained in detail in chapter 8, in the context of the LFV

9. This scale factor turns out to be close to 1.

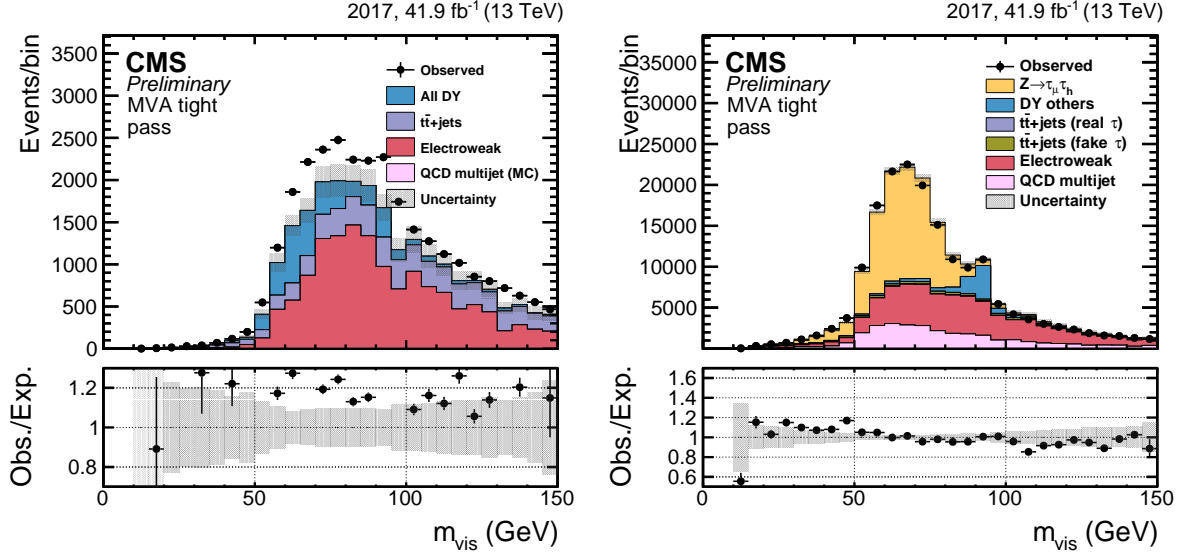


Figure 5.3 – Visible mass of the $\mu\tau_h$ system, in the W+jets control region (left, $m_T > 80$ GeV, before rescaling) and in the signal region (right, $m_T < 50$ GeV), using data collected by CMS in 2017. The tight anti-jet MVA discriminator, with the 2017v2 training is used. The W+jets background is lumped together with the diboson background and called “Electroweak”, but diboson events correspond to less than 1 % of W+jets events. The visible mass only takes into account the visible decay products of the tau, no attempt is made to recover the ν_τ . On the right plot, W+jets and QCD multijet events are estimated by a semi-data-driven procedure explained in-text.

analysis. The fake factor is allowed to depend on the τ_h p_T , decay mode, and its position in the barrel/endcap.

The signal region mass distribution obtained with this new method is shown in figure 5.4 (left). A fit still needs to be performed in order to determine the scale factor of the τ_h anti-jet discriminator efficiency, the resulting distribution is shown on the right side of figure 5.4. The nuisance parameters used for this fit are listed and explained below:

- Log-normal uncertainty¹⁰ on the luminosity: 2.5%. The integrated luminosity of a year of CMS data taking is known with finite precision. This uncertainty affects the yields of all processes estimated from simulations rather than from data-driven methods: the $Z \rightarrow \mu\tau_h$ signal, and the $Z \rightarrow \mu\mu$, diboson and $t\bar{t}$ backgrounds.
- Log-normal uncertainty on the $\mu \rightarrow \tau_h$ fake rate: 20%. A small number of $\mu\tau_h$ events are actually misidentified $\mu\mu$ events. The behavior of wrongly-identified particles is often poorly modeled by simulations, so a high uncertainty is warranted in this case.
- Log-normal uncertainty on the muon identification efficiency: 1% per muon (1% in the $\mu\tau_h$ final state, 2% in the $\mu\mu$ final state). This number is actually an uncertainty on the data/simulation muon identification scale factor estimated by the Muon POG. It affects all processes estimated from simulations: the $Z \rightarrow \mu\tau_h$ signal, and the $Z \rightarrow \mu\mu$, diboson and $t\bar{t}$ backgrounds.
- Log-normal uncertainty on the yield of the jet $\rightarrow \tau_h$ fakes background: 15%. This number is relatively large to account for the assumptions needed for the fake factor method to work (e.g. the fake factor is the same in the signal region and in the control region in which it was computed).

10. Cf. chapter 9 for an explanation of log-normal distributions.

- Log-normal uncertainty on the diboson and $t\bar{t}$ cross sections: 10%. These are conservative uncertainties, but overestimating them does not significantly affect the results since these are small backgrounds.
- Shape uncertainty on the tau energy scale: $\sim 1\%$, depending on the decay mode. This uncertainty affects all processes estimated from simulations, but is only applied to events in which the reconstructed τ_h matches a τ_h at the generated level of the simulation¹¹.
- Shape uncertainty on the energy scale of a muon misidentified as a τ_h : 3%. This is applied to all processes estimated from simulations.
- Shape uncertainty on the minimum-bias cross section: 4.6%. The minimum-bias cross section is the inclusive proton-proton cross section in LHC. Its value is used in the pileup reweighting procedure, as explained in chapter 6. This is applied to all processes estimated from simulations.
- Shape uncertainties for tau candidates which fail the anti-jet discriminator. Events with misidentified taus are estimated by applying a scale factor to data obtained in a control region close to the signal, but with tau candidates which fail the anti-jet discriminator. Not all of these events are $\text{jet} \rightarrow \tau_h$ fakes, some of them are genuine taus, and others are $\mu \rightarrow \tau_h$ fakes. These processes which do not involve $\text{jet} \rightarrow \tau_h$ fakes are estimated by simulation, and need to be subtracted from data in order to get the proper distribution of $\text{jet} \rightarrow \tau_h$ fakes. However, the normalizations of the distributions of events failing the anti-jet discriminators but with genuine τ_h and $\mu \rightarrow \tau_h$ fakes may be poorly described by simulations, indirectly affecting the estimation of $\text{jet} \rightarrow \tau_h$ fakes. Yield uncertainties of respectively 25% and 10% are applied to such events with $\mu \rightarrow \tau_h$ fakes and genuine taus.
- Decay mode dependent shape uncertainties on the fake factors, estimated by taking the difference between fake factors depending only on the τ_h decay mode and fake factors depending on τ_h decay mode, p_T and position in the barrel/endcap.
- Bin by bin uncertainties, modeling statistical fluctuations in each bin of the distribution, for all processes.

Shape uncertainties are fed into the model by providing two extra distributions on top of the nominal one, with each of the extra distributions providing a one standard deviation variation in one direction or another. The fitting tool then extrapolates a probability distribution from that. Only variations in shape are considered, differences in normalization are ignored.

With the fit I obtained various data/simulation scale factors for the efficiencies of various working points of the tau anti-jet discriminators (cut-based and BDT working points, I did not study the DeepTau algorithm), with dependence on the τ_h decay mode, p_T and position in the barrel/endcap. The latest results I obtained for the tight BDT working point are given in table 5.4. The inclusive scale factor is ~ 0.91 with an error of 0.04-0.05. The dominant source of uncertainty is the normalization of the $\text{jet} \rightarrow \tau$ background. Scale factors for 1-prong taus are close to 1, while for other decay modes they are typically around 0.9. Note that scale factor errors for the inclusive p_T range are sometimes bigger than for each of the p_T bins, indicating a larger fit instability. The Tau POG continued to work on this measurement after I left, and the final numbers they provided were different from these ones.

11. More information on the different aspects of event simulation is given in chapter 6.

Tau DM	$ \eta $ range	p_T range (GeV)	Scale Factor	Uncertainty
inclusive	inclusive	inclusive	0.912	-0.047/+0.037
inclusive	inclusive	20-40	0.928	-0.025/+0.026
inclusive	inclusive	40-150	0.922	-0.029/+0.033
DM0	barrel	inclusive	0.981	-0.057/+0.055
DM0	barrel	20-40	0.992	-0.036/+0.041
DM0	barrel	40-150	1.075	-0.075/+0.072
DM0	endcap	inclusive	1.008	-0.076/+0.077
DM0	endcap	20-40	1.01	-0.037/+0.042
DM0	endcap	40-150	0.845	-0.078/+0.078
DM1	barrel	inclusive	0.866	-0.022/+0.024
DM1	barrel	20-40	0.876	-0.025/+0.039
DM1	barrel	40-150	0.933	-0.029/+0.033
DM1	endcap	inclusive	0.925	-0.036/+0.037
DM1	endcap	20-40	0.931	-0.045/+0.064
DM1	endcap	40-150	0.854	-0.067/+0.089
DM10	barrel	inclusive	0.892	-0.026/+0.029
DM10	barrel	20-40	0.891	-0.031/+0.055
DM10	barrel	40-150	0.890	-0.050/+0.053
DM10	endcap	inclusive	0.919	-0.065/+0.069
DM10	endcap	20-40	0.917	-0.052/+0.071
DM10	endcap	40-150	0.883	-0.086/+0.101

Table 5.4 – Data/simulation scale factors for the efficiency of the tight working of the BDT anti-jet discriminator, with the 2017v2 training. The scale factors are characterized by the tau decay mode, location in the barrel/endcap, and p_T range. Decay mode 0 (DM0) corresponds to 1-prong + 0 π^0 decay, DM1 to 1-prong + 1 π^0 , and DM10 to 3-prong + 0 π^0 .

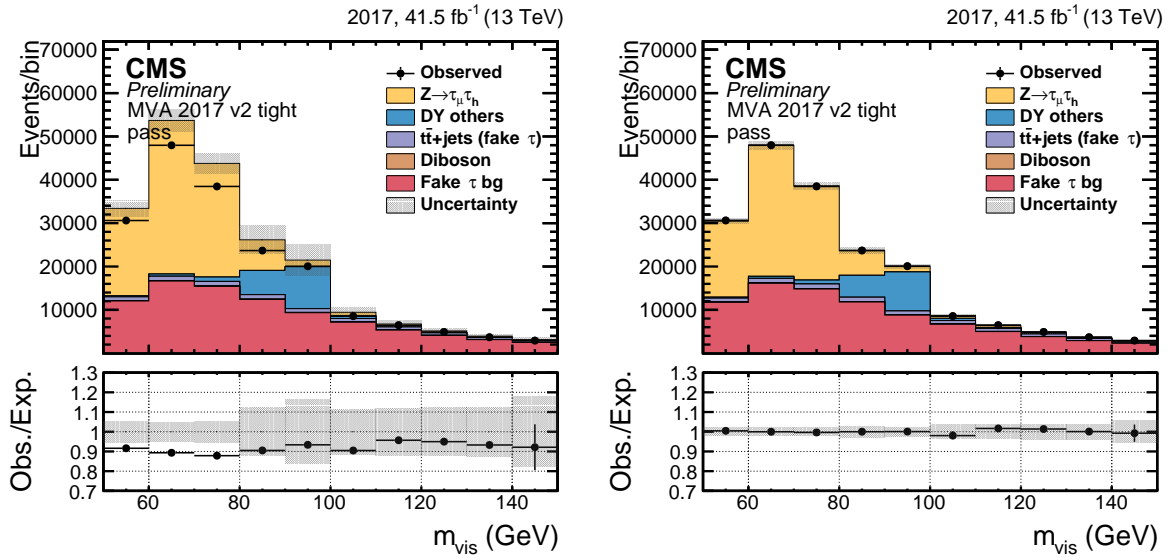


Figure 5.4 – Visible mass of the $\mu\tau_h$ system in the signal region, prefit (left) and postfit (right), using data collected by CMS in 2017. The tight anti-jet MVA discriminator, with the 2017v2 training is used. A wider binning compared to figure 5.3 was used in order to facilitate the fit convergence. The uncertainty bands are smaller postfit because the nuisance parameters are constrained by the fit to the data. The fit allows the extraction of the data/simulation scale factor for the tau anti-jet discriminator efficiency.

Chapter summary

The raw data collected by CMS needs to be reconstructed in a way that is usable by physics analyses. Tracker information is used to reconstruct tracks and interaction vertices, calorimeter information is used to build energy clusters, and muon system information is used to seed the reconstruction of muons. A global event description in CMS is provided by the Particle Flow (PF) framework, which correlates energy clusters and tracks in order to build particle candidates and classify them in mutually exclusive categories: muons, electrons, photons, charged hadrons, and neutral hadrons. Taus are reconstructed from charged hadrons and ECAL “strips”. Closely packed groups of particles are clustered into jets in a procedure meant to recover the properties of the partons which originate the jets. A missing transverse momentum vector is also calculated, providing a better description of events with invisible particles. However, particle identification criteria outside of the PF framework can also be used, because they sometimes allow for higher identification efficiencies of high p_T objects. This is the case for electron and muon identification in the LFV analysis. If non-PF objects are used, the missing energy needs to be corrected for the event description to remain consistent.

The performance of particle identification algorithms has to be checked in data. An example of such a performance study for the identification of taus was presented in detail, because I personally studied the efficiency of the cut-based and BDT-based tau anti-jet discriminators. I obtained efficiency scale factors for each of the algorithms, representing the ratio of the efficiency in data to the efficiency in simulations. The scale factors are extracted by first selecting events in a region enriched in $Z \rightarrow \tau_h \tau_\mu$ events, which are the signal events for this study. Then a fit of the background and signal events to the data is performed. The scale factors obtained were for the old tau anti-jet discriminators, and not the DeepTau algorithm.

My personal contributions

I was part of the Tau POG until May 2019 and studied the tau identification efficiency in 2017 data. I obtained efficiency scale factors for all the cut-based and BDT-based tau anti-jet discriminators, with the 2017v2 training. The scale factors were characterized by the discriminator working point, the tau decay mode (1-prong + 0 π^0 , 1-prong + 1 π^0 or 3-prong + 0 π^0), p_T range (20-40 or 40-150 GeV) and $|\eta|$ (barrel or endcap).

Part III

Analysis of lepton flavor violating final states

Datasets and simulations

Even after the CMS High Level Trigger has selected the most interesting proton-proton collision events for long term storage, the amount of data collected is enormous, corresponding to multiple petabytes each year [163]. Only a subset of these events is relevant for the data analyses performed in this thesis, and the specific datasets are detailed in section 6.1.

Event simulations are also an essential part of a data analysis: simulations can translate new physics models (e.g. lepton flavor violating - LFV new physics) into specific predictions for what should be observed in data, and they can also be used to estimate background events. Section 6.2 describes event simulation, and includes both a general overview of how it is performed in CMS and other collider experiments, and what specific simulations are used in this thesis for the LFV analysis.

6.1 Datasets

The data used in this thesis consists of proton-proton collision events at a center-of-mass energy $\sqrt{s} = 13 \text{ TeV}$, collected by CMS in the years 2016-18 during LHC Run 2, for a total integrated luminosity of 137.1 fb^{-1} . The luminosities for each yearly data-taking campaign are given in table 6.1. The LFV analysis uses data with at least one high momentum muon, electron or photon, which in 2016-17 corresponded to the CMS datasets called SingleMuon, SingleElectron and SinglePhoton. In 2018, the SingleMuon dataset stayed the same, but the SinglePhoton and SingleElectron datasets no longer existed and were replaced by the EGamma dataset, following a merger meant to increase data collection efficiency, as explained in chapter 4.

Data era	lumi (fb^{-1})
2016	35.9
2017	41.5
2018	59.7

Table 6.1 – Luminosities for each Run 2 data-taking campaign used in this thesis.

The LFV analysis is split into three channels, with the $e\mu$ channel using the SingleMuon and SinglePhoton datasets, the $e\tau$ channel using the SingleElectron and SinglePhoton datasets, and the $\mu\tau$ channel using the SingleMuon dataset. In 2018 the $e\mu$ and $e\tau$ channels used the EGamma dataset which replaced the SinglePhoton and/or SingleElectron datasets. The SingleMuon dataset is also used by the tau identification efficiency study presented in chapter 5.

Table 6.2 summarizes the purposes of each of the datasets used in this thesis.

Data-taking years	Datasets	LFV analysis			Tau ID
		$e\mu$	$e\tau$	$\mu\tau$	
2016-18	SingleMuon	✓		✓	✓
2016-17	SingleElectron	✓			
	SinglePhoton	✓	✓		
2018	EGamma	✓	✓		

Table 6.2 – CMS datasets used in this thesis, for the LFV analysis and the tau identification (tau ID) efficiency study.

6.2 Simulations

6.2.1 General overview of event simulation for collider experiments

Protons are composite particles, with three valence quarks, two ups and one down. But the picture gets more complicated when protons are probed at high enough energies: the valence quarks exchange gluons, which can themselves convert to quark-antiquark pairs, which can also exchange gluons, etc. As a result, high-energy proton-proton collisions are sometimes messy, with one specific parton (quark or gluon) from one proton hitting another parton from the other proton, creating one or more unstable particles and a flurry of hadrons. Those many particles then fly out into the detector, interact with it, and produce the electric signals which are eventually used by the reconstruction software to build a picture of the event that is as complete and accurate as possible. Event simulation programs deal with this complexity by dividing it into a series of simpler tasks which are mostly independent from each other [164]:

- **Hard interaction.** The physics processes which are relevant to this thesis, e.g. the production of vector bosons which subsequently decay into high-momentum leptons, are ultimately the result of an interaction between two partons, one from each of the colliding protons. The hard process is then simply the interaction of the two partons, which can be gluons or quarks of any flavor, and the production of the final state high-momentum leptons. The cross sections of hard processes can be calculated by perturbation theory, using Feynman diagrams such as the one in figure 2.1 (in chapter 2).
- **Partonic showers.** There are partons in the initial state of the hard process, and there are often partons in the final state too. They can radiate gluons, which can themselves radiate more gluons or create quark-antiquark pairs, in chain reactions called partonic showers. The momenta of individual partons in each shower start at the energy scale of the hard process and become lower as the shower progresses, eventually reaching an energy scale in which QCD interactions can no longer be calculated by perturbation theory. At this point, non-perturbative hadronization models take over.
- **Underlying event.** The partons involved in the hard process are not the only partons interacting in the event. These extra interactions produce additional partons, which are typically soft, and modify certain final state observables.
- **Hadronization.** For momentum scales below 1 GeV, non-perturbative models are needed to describe QCD interactions, and the confinement of partons into hadrons. The hadrons must be color-neutral, which means that partons cannot hadronize independently of each other, color-connected partons must instead hadronize together in order to ensure color-balancing. The hadronization models are phenomenological and mostly independent of

how the parton system was produced, which means that they can be tuned on specific datasets and then used to make predictions for other data.

- **Decays of unstable particles.** Some particles have a lifetime so short that they decay within the detector, and their decays must also be simulated.
- **Interaction with the detector.** The final state particles are propagated through a virtual model of the CMS detector, and their interactions with the material (tracker, calorimeters, muon chambers, solenoid magnet, service cables, cooling system, etc.) are simulated, as well as the detector response. The software used to simulate these detector interactions is Geant4 [165, 166].

The properties of simulated particles are saved at both the generated level (before their interaction with the detector is modeled) and the reconstructed level (using only the detector response).

The event cross section depends on how initial state partons are distributed within each of the colliding protons, which is why parton distribution functions (PDFs) [167] are defined. PDFs depend both on the energy scale (squared) Q^2 of the hard interaction and on the proton longitudinal momentum fraction x carried by the parton. The proton momentum fraction is not Lorentz invariant, but it is well defined in a boosted reference frame where the proton mass is negligible compared to its momentum. CMS is certainly such a reference frame for the 6.5 TeV protons that collide there.

A sample PDF function is the gluon PDF $f_g(x, Q^2)$, which if evaluated at $f_g(x = 0.01, Q^2 = 100 \text{ GeV}^2)$ is the probability density for finding a gluon with 1 % of the proton momentum, when the gluon is probed at an energy scale of 10 GeV.

There are many possible variations in the way events are simulated. Different softwares can be used, the hard interaction can be simulated at leading order (LO) or at higher orders of perturbation theory, the underlying event tune can change, and different showering and hadronization models can be used. Some extra details on the specific simulations used in this thesis are given in the next two sections.

6.2.2 Simulation of lepton flavor violating new physics

The simulated signal samples for the LFV analysis are the production and subsequent decay to LFV final states of RPV SUSY tau sneutrinos ($\tilde{\nu}_\tau$), quantum black holes (QBH), and Z' bosons. The Z' boson is assumed to have the same quark sector couplings as the SM Z boson, but in the lepton sector it decays exclusively to LFV final states. A dedicated sample is used for each of the three LFV final states. The relative width of the Z' boson is taken to be 3% of its mass, and interference effects with the SM are ignored. The QBHs are assumed to be spin-zero, colorless, electrically neutral, and to come from an ADD model with four extra spatial dimensions (cf. chapter 2). The RPV SUSY model is simplified, with all RPV couplings assumed to vanish except for the λ' couplings linking two same-flavor down-type quarks and the $\tilde{\nu}_\tau$ and one λ coupling allowing the $\tilde{\nu}_\tau$ to decay to a specific LFV final state ($e\mu$, $e\tau$ or $\mu\tau$). This means that a different λ coupling has a non-zero value for the RPV samples corresponding to each of the three LFV final states. All non-zero λ' and λ couplings are assumed to be equal. The yields of RPV samples are normalized to a next-to-leading order (NLO) calculation of the cross section.

For all three signals, events were generated at leading order, and the showering and hadronization simulations were performed by PYTHIA [168, 169, 170]. The underlying event tunes were determined using both CMS and CDF data [171], tune CUETP8M1 was used for 2016 samples

and tune CP5 for 2017-18 samples. Different generators were used for each of the signal samples: CalcHEP [172] for RPV events, version v3.0 of a dedicated QBH event generator [173] for QBH events, and PYTHIA8 for Z' events. The Z' boson samples were generated with NNPDF v3.1 sets [174] in 2017-18. The 2016 Z' samples were generated with NNPDF 3.0 sets [167] and are used in the $e\mu$ channel. However, they are not compatible with the use of the DeepTau algorithm and are thus not used in the tau channels (the 2017 samples are used instead for the 2016 analysis). RPV and QBH samples used CTEQ6L PDF sets [175].

Events were generated for different mass hypotheses of the Z' and $\tilde{\nu}_\tau$ resonances. For the RPV signal, a few different values of the LFV coupling constants were also simulated. Quantum black holes are not resonances, but different QBH models can be characterized by their threshold masses above which QBHs may start to be created. This threshold mass is related to the fundamental Planck mass, which for models with extra spatial dimensions can be arbitrarily low (e.g. ~ 1 TeV, as described in chapter 2). Tables 6.3-6.5 show the signal cross sections for some Z' , QBH, and RPV samples.

Z' Mass (GeV)	$\sigma (q\bar{q} \rightarrow Z' \rightarrow e\mu, e\tau, \mu\tau)$ (pb)
500	9.56
600	5.03
700	2.83
800	1.704
900	1.075
1000	0.7141
1100	0.4775
1200	0.329
1300	0.234
1400	0.1675
1500	0.1226
1600	9.071×10^{-2}
1700	6.808×10^{-2}
1800	5.166×10^{-2}
1900	3.912×10^{-2}
2000	3.027×10^{-2}
2200	1.847×10^{-2}
2400	1.147×10^{-2}
2600	7.258×10^{-3}
2800	4.695×10^{-3}
3000	3.079×10^{-3}
3500	1.163×10^{-3}
4000	4.84×10^{-4}
4500	2.2×10^{-4}
5000	1.11×10^{-4}

Table 6.3 – Cross sections of Z' samples as a function of the Z' resonance mass. The cross sections are assumed to be the same for all three LFV final states.

QBH Mass (GeV)	$\sigma(q\bar{q} \rightarrow \text{QBH} \rightarrow e\mu, e\tau, \mu\tau)$ (pb)
200	24066
400	1773
1000	32.8
2000	0.7475
5000	2.758×10^{-4}
10000	1.6×10^{-10}

Table 6.4 – Cross sections of ADD quantum black hole with 4 extra dimentions as a function of the QBH threshold mass. The cross sections are assumed to be the same for all three LFV final states.

$\tilde{\nu}_\tau$ Mass (GeV)	$\lambda = \lambda'$	$\sigma(q\bar{q} \rightarrow \tilde{\nu}_\tau \rightarrow e\mu, e\tau, \mu\tau)$ (pb)
200	0.01	773.72
300	0.01	212.55
400	0.01	81.57
500	0.01	37.73
600	0.01	19.66
700	0.01	11.130
800	0.01	6.6938
900	0.01	4.2154
1000	0.01	2.7521
1200	0.01	1.27272
1400	0.01	0.63690
1600	0.01	0.3369363
1800	0.01	0.185585
2000	0.01	0.105317
2500	0.01	2.7834×10^{-2}
3000	0.01	7.879×10^{-3}
3500	0.01	2.2935×10^{-3}
3500	0.1	2.2935×10^{-1}
4000	0.01	6.7061×10^{-4}
4000	0.1	6.704×10^{-2}
4000	0.2	2.682463×10^{-1}
4500	0.01	1.9445×10^{-4}
5000	0.01	5.5668×10^{-5}

Table 6.5 – Cross sections of RPV samples as a function of the $\tilde{\nu}_\tau$ resonance mass and the values of the λ and λ' coupling constants. $\lambda = \lambda'$ is assumed. The cross sections are assumed to be the same for all three LFV final states.

6.2.3 Simulation of Standard Model processes

The analysis selections, detailed in chapter 7, are optimized for the discovery of LFV new physics. Despite that, some background events from Standard Model (SM) processes still pass the selections, either because their final state contains two genuine different-flavor leptons, or because one or more of the leptons is actually a misidentified particle. In the $e\mu$ final state, these fake lepton backgrounds are mostly made up of events with jets faking electrons or muons, while in the tau channels a significant number of background events include jets faking tau leptons. The simulation of backgrounds with wrongly-identified leptons is unreliable, so these processes are estimated using data-driven “fake factor” methods, as explained in chapter 8. However, backgrounds with genuine leptons are estimated from simulations, as described in the following list:

- WW is one of the dominant analysis backgrounds. Each of the two W bosons can decay into a different lepton (+corresponding neutrino), mimicking an LFV final state. WW to dilepton samples, binned by the dilepton mass, were generated with POWHEG [176]. The cross sections are computed at the NLO level [177].
- $t\bar{t}$ is the other dominant background. Top quarks decay almost exclusively to a W boson and a bottom quark, so after a $t\bar{t}$ pair decays, there are typically two W bosons, which can generate LFV-like final states as in the WW process. Similarly to the WW samples, $t\bar{t}$ to dilepton samples were generated with POWHEG, in dilepton mass bins. The $t\bar{t}$ cross section is determined at the NNLO QCD + NLO electroweak level [178].
- Diboson backgrounds other than WW (WZ and ZZ) are also important. The samples were generated with MADGRAPH5_aMC@NLO [179, 180] and are classified according to their final states with at least two real leptons: WZ can yield 2 leptons + 2 quarks (2l2q) and 3 leptons + 1 neutrino (3l ν), while ZZ can yield 2l2q, 2l2 ν and 4l. The ZZ cross sections are computed at NNLO [181], while WZ cross sections are NLO [182].
- DY to dilepton events, with any number of accompanying jets (DY+jets), are a small background. If a ditau pair is produced, then many LFV-like final states are possible, as the taus can decay into electrons, muons or hadrons. Other possibilities include the wrong identification of one of the lepton flavors, or a jet faking an electron or a tau. Like the WZ and ZZ processes, the DY+jets background was generated at NLO precision in dilepton invariant mass bins with MADGRAPH5_aMC@NLO. The cross sections are computed at the NLO level as well [183].
- Events with a single top quark also contribute to the background. Separate samples for events with top and antitop quarks were generated with POWHEG. Both cross sections are identical, and are computed at the NNLO level [184].

The hadronization and parton showering were modeled with PYTHIA and the underlying event tune was CUETPM1 in 2016 and CP5 in 2017-18, same as for the signal samples. The NNPDF v3.0 PDF [167] sets were used in 2016, while the v3.1 sets [174] were used in 2017-18.

Table 6.6 shows the simulated background processes used in the LFV analysis, and their respective cross sections. Note that some of these processes are also useful to the tau ID study presented in chapter 5. Almost all processes in the table have two real leptons in the final state, with some exceptions: in some WZ and ZZ samples the W and Z bosons are not required to decay leptonically, the jets in the DY+jets samples can fake leptons, and W+jets events can only be selected when a jet fakes a lepton. Nevertheless, simulated W+jets events are useful in order to check the robustness of the fake factor method in the tau channels. As for the other simulated events with jet \rightarrow lepton fakes, they are filtered out by the analysis code using generated level information, because any process with jet \rightarrow lepton fakes is exclusively estimated by data-driven methods.

6.2.4 Normalization, pileup reweighting and other scale factors

Simulated distributions need to be normalized to the data integrated luminosity, according to the cross section of each process. The following formula is used:

$$f^{\text{norm}} = f^{\text{sim}} \cdot \frac{\sigma L}{N_{\text{events}}} \quad (6.1)$$

where f^{norm} and f^{sim} are respectively the normalized and non-normalized distributions, σ is the cross section of the process, L is the integrated luminosity of the data and N_{events} is the total number of events in the simulated sample. f can be any distribution, e.g. for the $\mu\tau$ final

Process	mass/ p_T bins	Cross section (pb)	LFV analysis	Tau ID
DY($\rightarrow ll$)+jets	50 GeV $< m_{ll}$	6077.22	✓	✓
	100 GeV $< m_{ll} < 200$ GeV	226.6	✓	
	200 GeV $< m_{ll} < 400$ GeV	7.77	✓	
	400 GeV $< m_{ll} < 500$ GeV	0.4259	✓	
	500 GeV $< m_{ll} < 700$ GeV	0.2446	✓	
	700 GeV $< m_{ll} < 800$ GeV	0.0378	✓	
	800 GeV $< m_{ll} < 1\,000$ GeV	0.0318	✓	
	1\,000 GeV $< m_{ll} < 1\,500$ GeV	0.0202	✓	
	1\,500 GeV $< m_{ll} < 2\,000$ GeV	0.002286	✓	
	2\,000 GeV $< m_{ll} < 3\,000$ GeV	0.0005375	✓	
$t\bar{t} \rightarrow 2l2\nu$	inclusive	87.33	✓	✓
	500 GeV $< m_{ll} < 800$ GeV	0.326	✓	
	800 GeV $< m_{ll} < 1\,200$ GeV	0.0326	✓	
	1\,200 GeV $< m_{ll} < 1\,800$ GeV	0.00305	✓	
	1\,800 GeV $< m_{ll}$	0.000174	✓	
WW $\rightarrow 2l2\nu$	inclusive	12.178	✓	✓
	200 GeV $< m_{ll} < 600$ GeV	1.39	✓	
	600 GeV $< m_{ll} < 1\,200$ GeV	0.057	✓	
	1\,200 GeV $< m_{ll} < 2\,500$ GeV	0.0036	✓	
	2\,500 GeV $< m_{ll}$	0.000054	✓	
WZ $\rightarrow 2l2q$		5.595	✓	✓
WZ $\rightarrow 3l\nu$		5.052	✓	✓
ZZ $\rightarrow 2l2\nu$		0.564	✓	✓
ZZ $\rightarrow 2l2q$		3.22	✓	✓
ZZ $\rightarrow 4l$		1.212	✓	✓
Single top (top)		19.47	✓	
Single top (antitop)		19.47	✓	
W($\rightarrow l\nu$)+jets	inclusive	60430		
	50 GeV $< p_T(W) < 100$ GeV	3046		
	100 GeV $< p_T(W) < 250$ GeV	627.1		
	250 GeV $< p_T(W) < 400$ GeV	21.83		
	400 GeV $< p_T(W) < 600$ GeV	2.635		
	600 GeV $< p_T(W)$	0.4102		

Table 6.6 – List of background samples and their cross sections (in pb). When two separate samples overlap (e.g. when there are both a mass-inclusive and mass-binned samples for the same process), the overlap is removed by checking generator-level information. The W+jets samples are only used to validate the data-driven method to estimate backgrounds with jet $\rightarrow \tau$ fakes. The W+jets samples were generated with MADGRAPH5_aMC@NLO. 2018 mass-binned $t\bar{t} \rightarrow 2l2\nu$ and WW $\rightarrow 2l2\nu$ samples are missing, so the 2017 samples are used instead. Some 2018 W+jets p_T bins are missing as well.

state the number of events binned by the p_T of the tau lepton.

However, the experimental conditions and the detector response cannot be perfectly modeled, and further reweighting of simulations may be needed to improve their agreement with data. The LFV analysis reweights the samples with CMS recommended scale factors, measured by the relevant particle object groups (POGs): muon trigger (for the $e\mu$ and $\mu\tau$ channels), electron trigger (for the $e\tau$ channel), muon reconstruction and isolation, electron identification, tau reconstruction and identification, $e \rightarrow \tau$ and $\mu \rightarrow \tau$ misidentification (the misidentification is checked using generated level information), tracking efficiency (for 2016 only). The energy scale of reconstructed taus also needs to be corrected, depending on the tau decay mode and whether the reconstructed tau is genuine or a muon or electron fake, as explained in chapter 5. This energy scale correction and the scale factors may change the overall normalization and the shape of the samples.

One final reweighting of the samples is necessary: pileup reweighting. At typical LHC instantaneous luminosities, there is more than one proton-proton interaction at each bunch crossing. The “interesting” hard interaction selected by the analysis trigger is surrounded by many other interactions, which are overwhelmingly likely to be soft QCD multijet collisions called pileup. These pileup interactions are simulated since many variables depend on the number of pileup interactions, such as the PF MET and the tau isolation. So-called minimum bias events, representing the average LHC collision, are used to simulate the pileup interactions. The number of such interactions is a function of the LHC instantaneous luminosity and thus changes with time. As the pileup distributions are not identical in data and in the simulations, the latter need to be reweighted for their shapes to show a better agreement with what is observed in data. This reweighting does not change the overall normalization of the simulations. The pileup profile is not estimated from the number of reconstructed vertices, since the efficiency is less than 100% and could differ between data and simulations. Instead, it is calculated from the instantaneous luminosity and the total inelastic proton-proton cross section of 69.2 mb.

Chapter summary

The LFV analysis uses LHC Run 2 proton-proton collisions data, collected in 2016-18 with the CMS detector. The center-of-mass energy of the collisions is 13 TeV, and the total integrated luminosity of the data is 137.1 fb^{-1} . Datasets with a single muon, single electron or single photon are used. For both the LFV analysis and the tau ID efficiency study, simulations of SM processes with two real leptons are used. For the LFV analysis specifically, signal samples of RPV SUSY, QBH and Z' processes were also generated. All aspects of a collision event are simulated: the parton distributions within each proton, the hard interaction between two of the partons, the parton showering, the underlying event, the hadronization, the decays of unstable particles, and the interaction of final state particles with the detector and the detector response. The simulations then need to be normalized according to their cross section and the integrated luminosity of the data. Extra scale factors are also applied to correct for the detector response measured in data. The tau energy scale is also corrected. Finally, the samples are reweighted according to their pileup for the shape of their pileup distributions to better conform to the data distribution.

Analysis strategy

This chapter summarizes how the search for lepton flavor violating (LFV) new physics at high invariant masses¹ is organized. Section 7.1 details the triggers and event selections applied in each of the three LFV final states, section 7.2 explains which discriminating mass or mass-like variables are used to distinguish the signal from the background, section 7.3 presents some general comments on the background estimation strategy with a particular emphasis on the $t\bar{t}$ process, and section 7.4 closes the chapter with a discussion of why the DeepTau algorithms are preferred to the old BDT algorithms for tau identification.

7.1 Trigger paths and event selection

The goal of the analysis is to search for lepton flavor violation in the charged lepton sector at the energy frontier. The relevant event signature is one in which there are two different-flavor leptons with high momenta. Trigger paths requiring the presence of (at least) one high-momentum lepton are thus ideally suited for the analysis, as they are simple, unprescaled and unlikely to veto signal events. Single muon triggers are used whenever possible (in the $e\mu$ and $\mu\tau$ channels) because muons are the particles CMS can most reliably reconstruct. For the $e\tau$ channel, single electron triggers are used because the tau triggers are less efficient and would be incompatible with the way we estimate $\text{jet} \rightarrow \tau$ fakes. The choice of triggers also explains the choice of datasets detailed in chapter 6, since each trigger is assigned to a specific dataset. The specific triggers used vary for each of the data-taking years for various reasons: sometimes old triggers become unavailable, or p_T and E_T thresholds are raised in order to keep the trigger rate under control as instantaneous luminosity increases. Table 7.1 summarizes which triggers are used in which channel and which year.

	2016	2017	2018
$e\mu$	Mu50 OR TkMu50 OR Photon175	Mu50 OR TkMu100 OR OldMu100 OR Photon175	Mu50 OR TkMu100 OR OldMu100 OR Photon200
$e\tau$	Ele27_WPTight OR Photon175 OR Ele115_CaloIdVT_GsfTrkIdT	Ele35_WPTight OR Photon175 OR Ele115_CaloIdVT_GsfTrkIdT	Ele32_WPTight OR Photon200 OR Ele115_CaloIdVT_GsfTrkIdT
$\mu\tau$	Mu50 OR TkMu50	Mu50 OR TkMu100 OR OldMu100	

Table 7.1 – Triggers used in each of the three LFV channels and each of the three data-taking years.

1. See the invariant mass definition in appendix A.

For the $e\mu$ channel, a photon trigger (Photon175 in 2016-17, Photon200 in 2018) is used in addition to the single muon triggers in order to compensate for an efficiency loss at high invariant mass $m_{e\mu}$. The photon trigger selects photon candidates with $E_T > 175$ GeV or 200 GeV but in practice it selects almost 100% of high E_T electrons as well, as shown in figure 7.1. The 2016 single muon triggers are Mu50 OR TkMu50, selecting muons with $p_T > 50$ GeV by two different algorithms. In later years, the TkMu50 is no longer available, and the Mu50 OR TkMu100 OR OldMu100 combination is used for the years 2017-18.

For the $e\tau$ channel, single electron and single photon triggers are used. The single electron triggers are a combination of a trigger with a lower E_T threshold but tighter identification requirements and another trigger with a higher E_T threshold but looser identification requirements. The E_T threshold of the first trigger varied from year to year: Ele27_WPTight in 2016, Ele35_WPTight in 2017 and Ele32_WPTight in 2018. The second electron trigger remained the same throughout the years: Ele115_CaloIdVT_GsfTrkIdT. Same as for the $e\mu$ channel, the photon trigger increases efficiency at high invariant masses, by selecting electrons missed by the electron triggers. Photon175 is used for 2016-17 data, and Photon200 for 2018 data.

For the $\mu\tau$ channel, only single muon triggers are used. Photon triggers would not measurably increase efficiency because there is a tau rather than an electron in the final state. And as already mentioned, tau triggers would be incompatible with the way we estimate events with $\text{jet} \rightarrow \tau$ fakes, so there are no good ways to improve trigger efficiency in the $\mu\tau$ channel when using Run 2 data². The single muon triggers are the same as those used in the $e\mu$ channel.

On top of the trigger requirements, the offline selections detailed in table 7.2 are applied. Since the number of SM events in LFV final states is already expected to be very low, selections are kept to a minimum in order to make the search as model-independent as possible. The two final state leptons are required to be well separated, with $\Delta R(e, \mu) > 0.1$, $\Delta R(e, \tau) > 0.5$ and $\Delta R(\mu, \tau) > 0.5$. η requirements guarantee that the reconstructed particles are within the detector acceptance, and minimum thresholds on the muon p_T or electron E_T ensure that the trigger is maximally efficient for selected events. As shown in figure 7.2, the Mu50 trigger is not immediately maximally efficient for muons above the 50 GeV p_T threshold, due to differences in the way the muon p_T is calculated. Muons are further required to pass the high- p_T identification and to have a tracker-based isolation < 0.1 , electrons have to pass the HEEP ID, and taus need to pass the decay mode finding algorithm, the tight anti-jet, the loose anti-electron and the tight anti-muon discriminators from the DeepTau algorithm. The minimum p_T threshold for the taus is 50 GeV because lower p_T tau candidates caused problems for the $\text{jet} \rightarrow \tau$ fakes estimation, especially for 2016 data. All channels are required to pass the MET filters since events with abnormally large MET should be excluded from the data analysis. Extra lepton vetos are applied to avoid any overlap between the three LFV channels: events with high- p_T muons are vetoed in the $e\tau$ channel and events with HEEP electrons are vetoed in the $\mu\tau$ channel. Extra lepton vetos are also applied to suppress events with same flavor leptons: events with two or more electrons are vetoed in the $e\tau$ channel, and events with a well separated dimuon pair are vetoed in the $\mu\tau$ channel. One final selection is applied in the tau channels in order to reduce the number of events with $\text{jet} \rightarrow \tau$ fakes: the transverse mass m_T of the light lepton and MET needs to be above 120 GeV. m_T is defined in the following way:

$$m_T = \sqrt{2 |\vec{p}_T^l| |\vec{p}_T| (1 - \cos \Delta\phi(\vec{p}_T^l, \vec{p}_T))} \quad (7.1)$$

where ϕ is the azimuthal angle, \vec{p}_T^l is the transverse momentum vector of the light lepton (electron or muon, for respectively the $e\tau$ and $\mu\tau$ final states) and \vec{p}_T is the MET vector. If

2. Improvements in trigger efficiency could be made in future data-taking runs.




Figure 7.1 – $e\mu$ channel. Trigger efficiencies of different triggers and trigger combinations in mass range $[0, 1000]$ GeV (top) and $[0, 7000]$ GeV (bottom), estimated using 2017 $Z' \rightarrow e\mu$ samples. The photon trigger has close to 100% efficiency at high invariant mass, allowing a ~ 10 -20% drop in muon trigger efficiency to be recovered.

a W boson decays in the transverse plane to a light lepton and neutrino, if the light lepton mass is approximated to be zero, and if there are no other sources of MET in the event, then the m_T formula recovers the mass of the W boson. W+jets events³ are thus expected to have m_T values close to $m_W = 80$ GeV, which is why they are suppressed by requiring $m_T > 120$ GeV.




Figure 7.2 – Mu50 trigger efficiency in part of 2018 data, as a function of offline $p_T(\mu)$ [185]. (Two histograms are plotted, before and after an intervention meant to increase efficiency, but this is not important here.) The trigger does not attain maximum efficiency when $p_T(\mu) = 50$ GeV, instead it does so shortly afterwards, because of differences between the p_T measurement at HLT and offline. This “turn-on curve” is observed in other triggers too.

A veto on b-jets, using the medium working of the CSV algorithm (the CSVv2 algorithm in 2016, DeepCSV in 2017-18), was considered in order to reject $t\bar{t}$ events, one of the main backgrounds of the analysis. Ultimately, it was decided not to use such a veto in order to keep the analysis as model-independent as possible. However, the selection of events with two b-jets is still useful in order to create a control region enriched in $t\bar{t}$ events and check how well they are estimated by the simulations. More details on this study are written in section 7.3.

7.2 Discriminating variables: invariant mass and collinear mass

The RPV SUSY $\tilde{\nu}_\tau$ and the Z' bosons are both hypothetical resonances, which means that they should show up in data as localized excesses above the SM background in the dilepton invariant mass spectrum. QBHs are not resonances, but there is still a threshold mass below which they are not created, and for higher masses their production cross section drops sharply, so the signal should still show up as a relatively localized excess in the mass spectrum.

The dilepton invariant mass is thus a useful discriminating variable between the signal and the background. In the $e\mu$ channel, there is no neutrino and thus no missing energy, and it is straightforward to calculate the invariant mass $m_{e\mu}$ using the reconstructed HEEP electron

3. W+jets events form a considerable fraction of events with $\text{jet} \rightarrow \tau$ fakes.

	$e\mu$	$e\tau$	$\mu\tau$
Muon	$p_T > 53 \text{ GeV}$, $ \eta < 2.4$ high pt muon ID, tracker-based isolation < 0.1		$p_T > 53 \text{ GeV}$, $ \eta < 2.4$ high pt muon ID, tracker-based isolation < 0.1
Electron	$E_T > 35 \text{ GeV}$, HEEP ID V7.0 (V7.0-2018Prompt for 2018), no μ with $p_T > 5 \text{ GeV}$ such that $\Delta R(e, \mu) < 0.1$	$p_T > 50 \text{ GeV}$, HEEP ID V7.0 (V7.0-2018Prompt for 2018)	
Tau		$p_T > 50 \text{ GeV}$, $ \eta < 2.3$, new DM finding (veto 2-prong tau: DM5,6), Deep τ tight anti-jet, loose anti-e, tight anti- μ discr.	$p_T > 50 \text{ GeV}$, $ \eta < 2.3$, new DM finding (veto 2-prong tau: DM5,6), Deep τ tight anti-jet, loose anti-e, tight anti- μ discr.
$m_T(l, \cancel{p}_T)$ MET filters ΔR	applied $\Delta R(e, \mu) > 0.1$	$> 120 \text{ GeV}$ applied $\Delta R(e, \tau) > 0.5$	$> 120 \text{ GeV}$ applied $\Delta R(\mu, \tau) > 0.5$
Extra lepton veto		Rejected events: events with (i) two or more electrons with $p_T > 10 \text{ GeV}$, $ \eta < 2.5$, VID cut-based ele ID veto (94X_V1 version) or (ii) a muon with $p_T > 35 \text{ GeV}$, $ \eta < 2.4$, HighPt ID, tracker isolation < 0.15	Rejected events: events with (i) a HEEP ID V7.0 electron with $p_T > 35 \text{ GeV}$ or (ii) a dimuon pair, with $\Delta R > 0.2$, and where both muons have: HighPt ID, $p_T > 10 \text{ GeV}$, $ \eta < 2.4$, tracker iso < 0.15 .

Table 7.2 – Selections used for each of the three channels in each of the data-taking years. The selection on the electron p_T/E_T is stricter in the $e\tau$ channel because an electron trigger is used, and a high threshold is necessary in order to avoid the turn-on curve. Details on the electron HEEP identification, the muon high- p_T identification, the MET filters, etc. are available in chapter 5.

and high- p_T muon. But in the tau channels, part of the final state energy is lost in the neutrino produced by tau decay. There are many possible mass variables that can be reconstructed. The simplest one is the visible mass m_{vis} which simply ignores the MET and computes the invariant mass using the light lepton and the visible components of the tau. Alternatively, one can take the MET into account by treating it as a 4-momentum vector confined to the transverse plane. The total mass m_{tot} is obtained by adding this MET 4-momentum to the light lepton and visible tau 4-momenta and calculating the invariant mass of the resulting total 4-momentum. A final, more sophisticated possibility is to take into account the kinematics of a signal event: the MET comes only from the neutrino resulting from the tau decay, and furthermore the high momentum of the tau implies that its decay products, neutrino included, are tightly collimated around a central axis. The invisible neutrino should then be approximately collinear to the visible decay products of the tau. By taking the part of the MET that is collinear with the τ_h p_T , calculating a visible fraction x_{vis} of the transverse momentum, and then correcting the τ_h 4-momentum by dividing it by x_{vis} , it is possible to calculate the collinear mass m_{col} :

$$\begin{aligned} m_{\text{col}} &= \frac{m_{\text{vis}}}{\sqrt{x_{\text{vis}}}} \\ x_{\text{vis}} &= \frac{p_T(\tau_h)}{p_T(\tau_h) + \cancel{p}_T \text{ col}} \\ \cancel{p}_T \text{ col} &= \max\left(\frac{\vec{\cancel{p}}_T \cdot \vec{p}_T(\tau_h)}{p_T(\tau_h)}, 0\right) \end{aligned} \quad (7.2)$$

where $\vec{\cancel{p}}_T$ is the MET vector, and $\vec{p}_T(\tau_h)$ and $p_T(\tau_h)$ are respectively the τ_h visible transverse momentum vector and the module of said vector. Note that dividing m_{vis} by $\sqrt{x_{\text{vis}}}$ is equivalent to dividing $p_T(\tau_h)$ by x_{vis} if the masses of the tau and the light lepton are neglected. Also note that when the MET is misaligned with the visible decay products of the tau, $\vec{\cancel{p}}_T \cdot \vec{p}_T(\tau_h)$ becomes negative. This happens in 2-7% of signal events, possibly when the MET is wrongly reconstructed. In these cases the collinear approximation becomes meaningless, so $\cancel{p}_T \text{ col}$ is set to zero, and $m_{\text{col}} = m_{\text{vis}}$.

For each signal, the three different mass variables are compared to each other in figures 7.3-7.5 for the $\mu\tau$ final state. 2016 samples are used for QBH and RPV, though no 2016 Z' samples compatible with the DeepTau ID are available. However, the results are similar for samples from different years. To aid the comparison, the mass resolution σ is defined by taking the normalized difference between the reconstructed mass variable m_{reco} and the generated mass m_{gen} :

$$\sigma(m_{\text{reco}}) = \frac{m_{\text{reco}} - m_{\text{gen}}}{m_{\text{gen}}} \quad (7.3)$$

m_{reco} better describes m_{gen} when the σ distribution is narrower, and its average is closer to zero. The means and the standard deviations of the σ distributions can be checked in table 7.3, for each signal process and each of the reconstructed mass variables. The figures and the table both show that m_{vis} is an inadequate variable, with a large standard deviation, and providing a large underestimation of m_{gen} . Different mass hypotheses for the signal show a sizeable overlap in their m_{vis} distributions, even when separated by 1 TeV or more. The m_{tot} variable is a considerable improvement, with much narrower and more symmetrical σ distributions and average values much closer to zero. However, the m_{col} variable is even better, which is why it was chosen as the discriminating analysis variable in the tau channels.

The m_{col} distributions in figure 7.5 also illustrate the fact that QBHs are not resonances. The distributions are asymmetrical, with few events below the threshold mass but a long tail

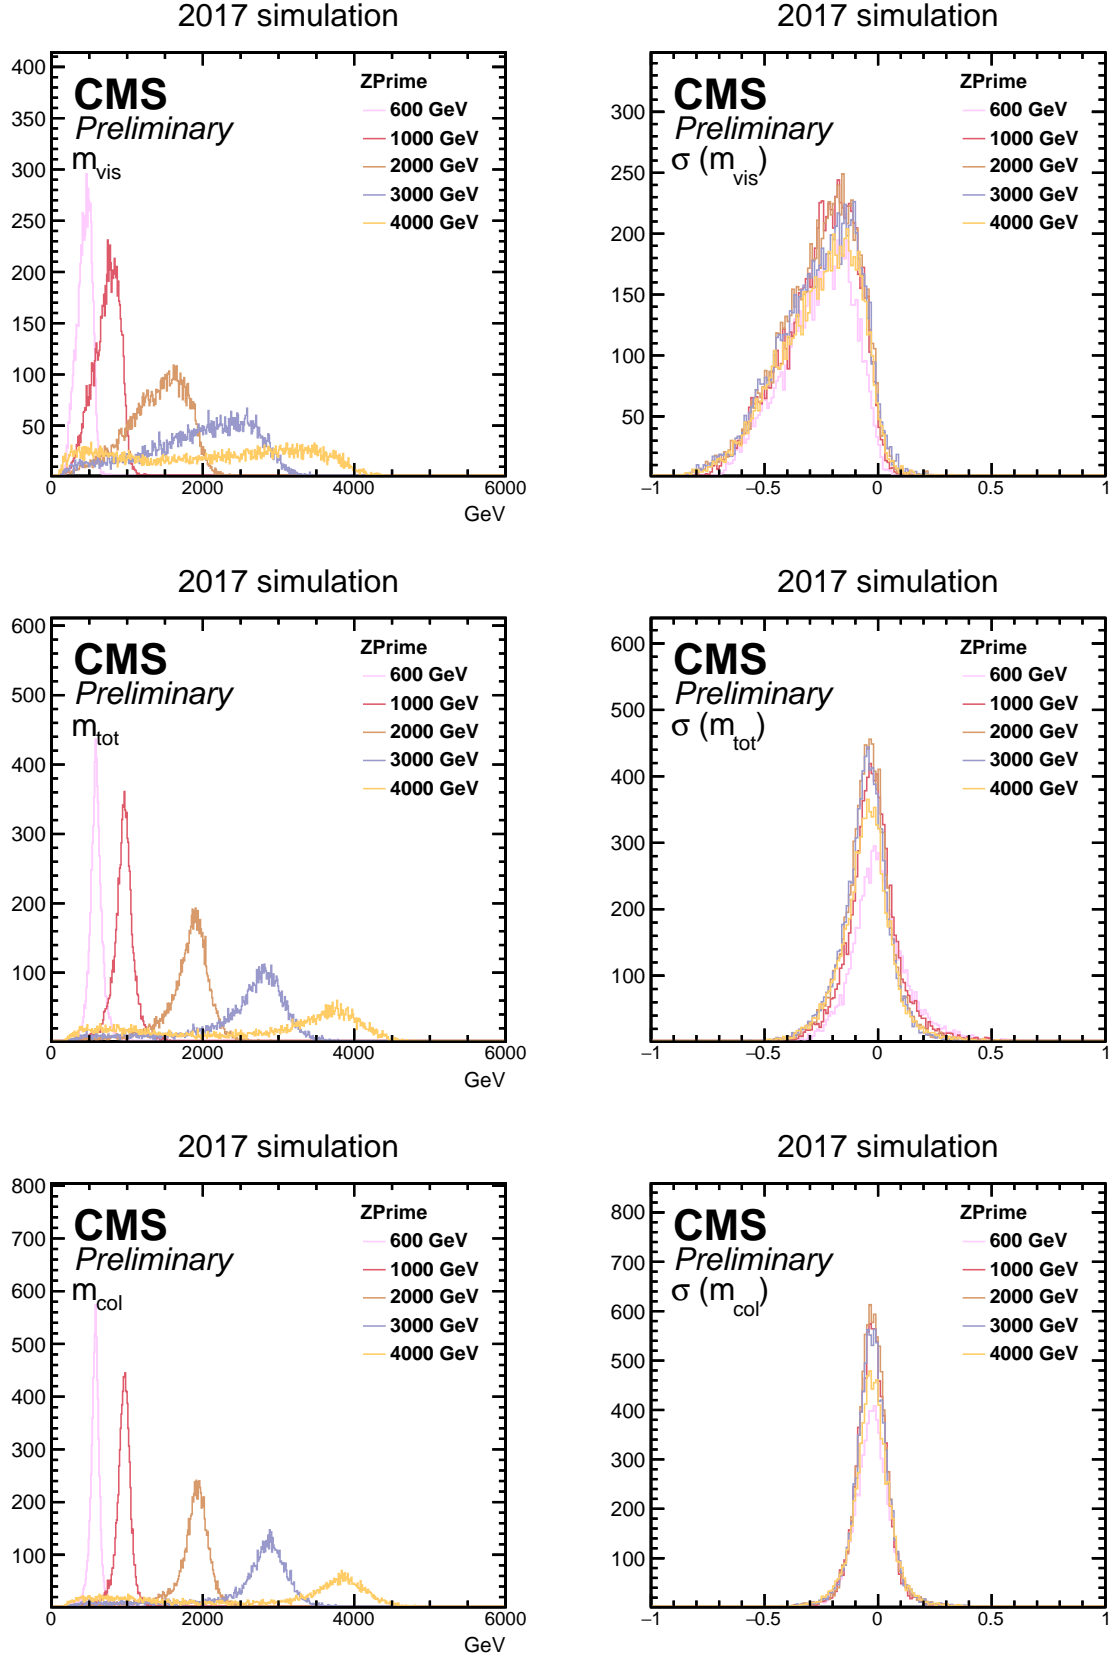


Figure 7.3 – Z' signal, comparison of the three reconstructed mass variables in the $\mu\tau$ final state: m_{vis} , m_{tot} and m_{col} . On the left the reconstructed mass distributions are shown, for various resonance mass hypotheses, normalized to the same (arbitrary) cross section. On the right, the mass resolution distributions are shown for the same mass points.

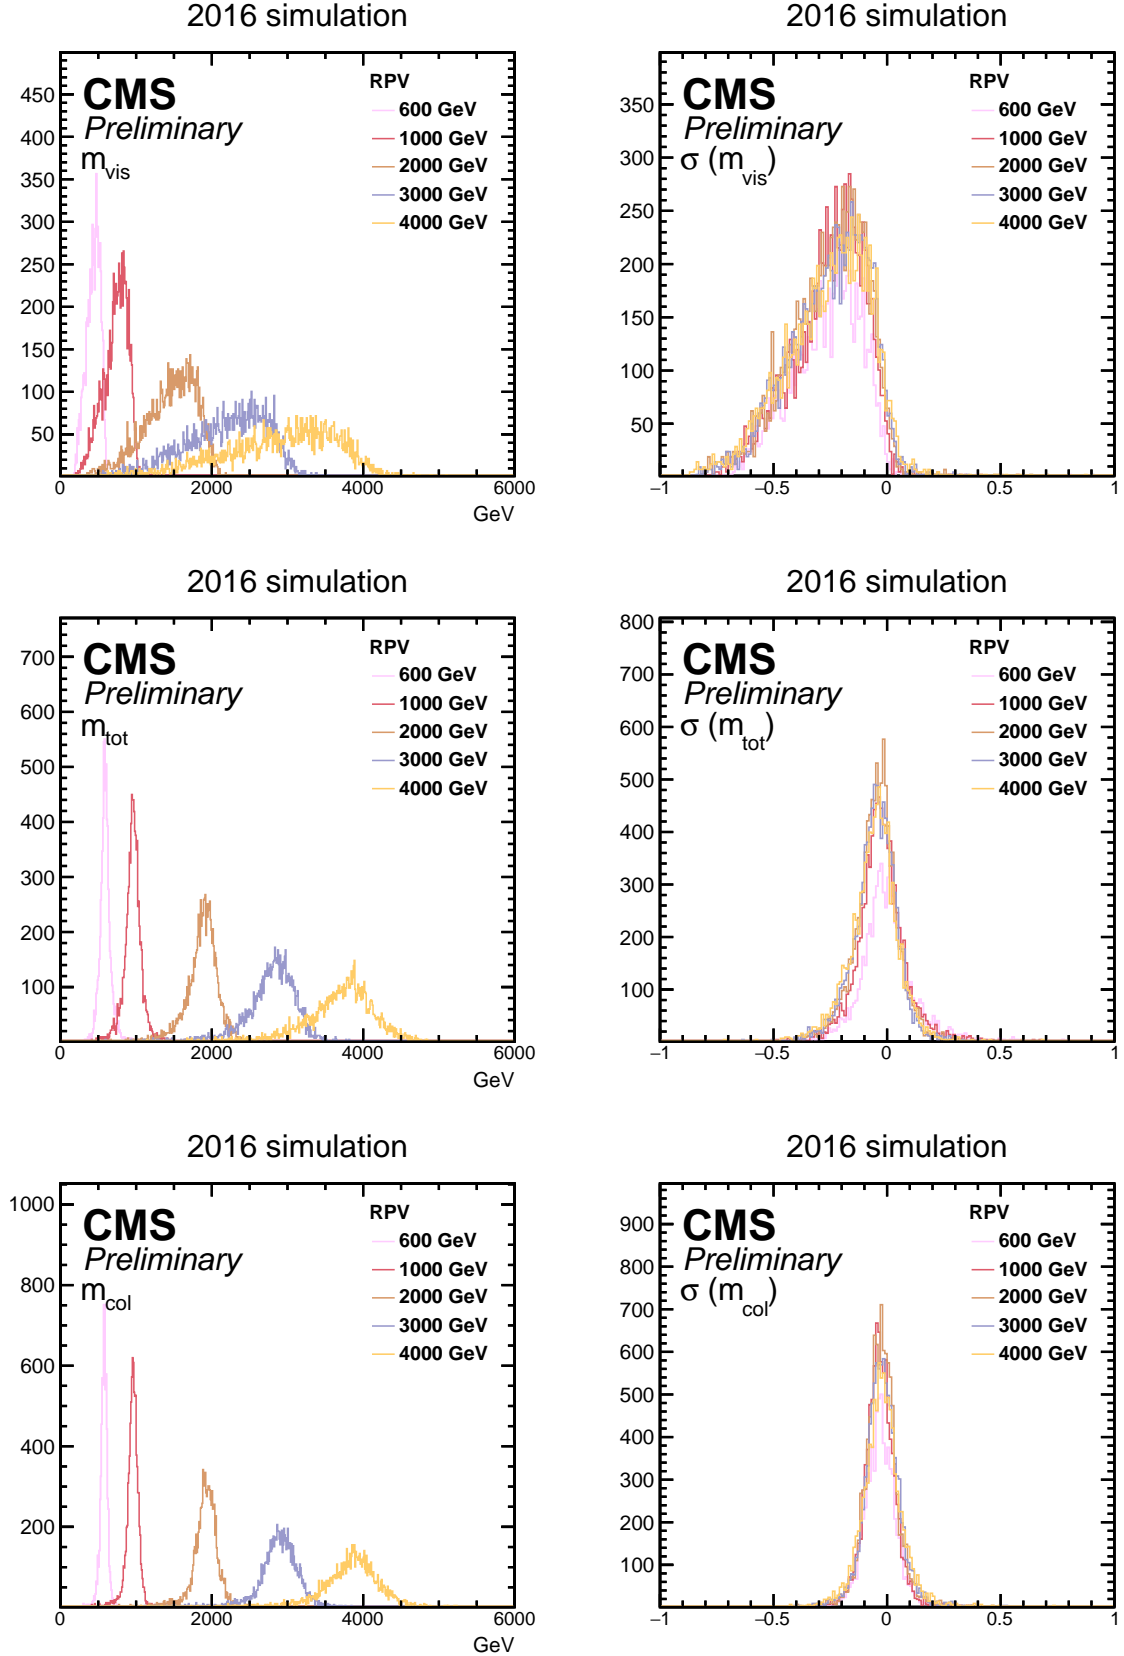


Figure 7.4 – RPV $\tilde{\nu}_\tau$ signal, comparison of the three reconstructed mass variables in the $\mu\tau$ final state: m_{vis} , m_{tot} and m_{col} . On the left the reconstructed mass distributions are shown, for various resonance mass hypotheses, normalized to the same (arbitrary) cross section. On the right, the mass resolution distributions are shown for the same mass points.

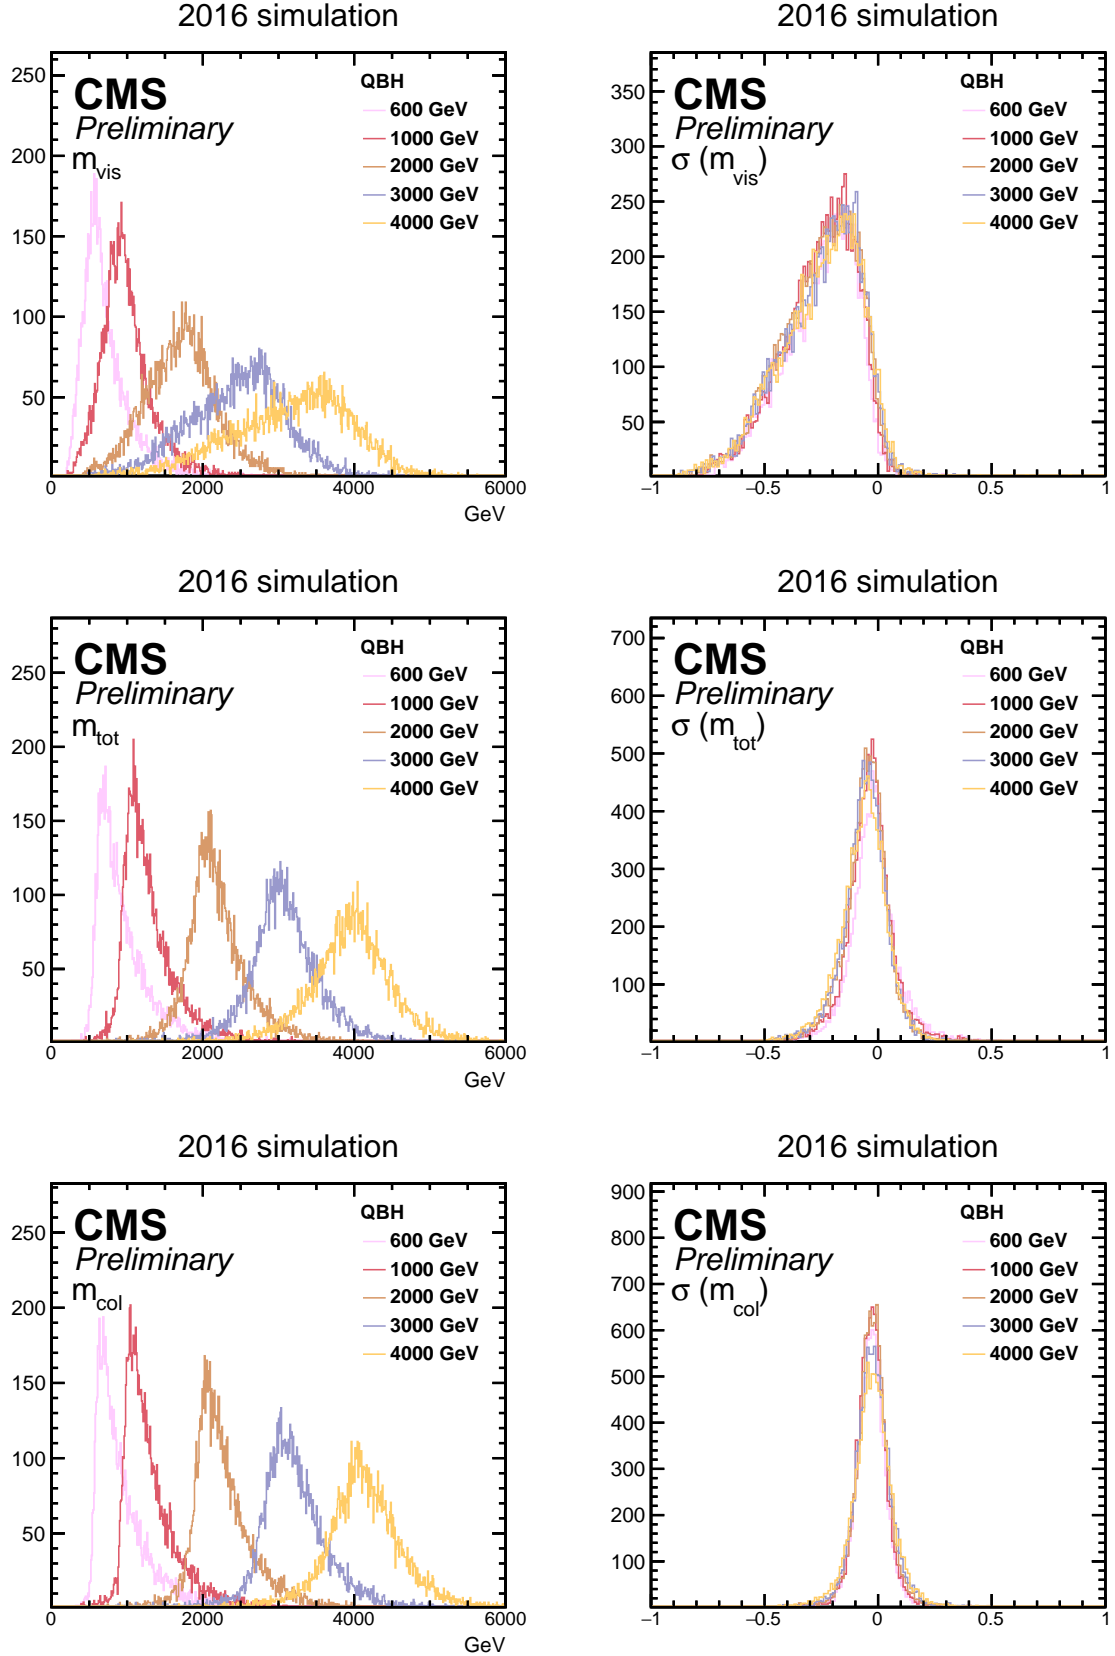


Figure 7.5 – QBH signal, comparison of the three reconstructed mass variables in the $\mu\tau$ final state: m_{vis} , m_{tot} and m_{col} . On the left the reconstructed mass distributions are shown, for various threshold mass hypotheses, normalized to the same (arbitrary) cross section. On the right, the mass resolution distributions are shown for the same mass points.

		m_{vis}	m_{tot}	m_{col}
Z' boson	mean	-0.26	-0.048	-0.026
	std. dev.	0.18	0.13	0.098
RPV SUSY $\tilde{\nu}_\tau$	mean	-0.26	-0.057	-0.022
	std. dev.	0.19	0.12	0.10
QBH	mean	-0.25	-0.064	-0.023
	std. dev.	0.18	0.11	0.10

Table 7.3 – Mean and standard deviation of the mass resolution distribution for each reconstructed mass variable and each of the three signal processes. The mass point shown for all three signals is 4 TeV, corresponding to the resonance mass for the $\tilde{\nu}_\tau$ and Z' boson, and the threshold mass for the QBH. m_{col} has the narrowest and least biased distributions.

above it. The cross section for the production and decay of a QBH is a decreasing continuum with a threshold: there is a minimum energy scale (close to the Planck mass) below which QBHs cannot be produced, and the production of QBHs with higher masses is disfavored due to the decreasing PDFs. This can be observed more clearly in the lowest mass points. Nevertheless, m_{col} remains a faithful mass reconstruction variable, as evidenced by the corresponding σ distributions.

As mentioned in the explanation of equation (7.2), in some signal events the MET is misaligned with the visible decay products of the tau, i.e. $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) > \frac{\pi}{2}$. This happens in both the $e\tau$ and the $\mu\tau$ channels. A study was performed in the $\mu\tau$ channel to determine how close these events are to normal signal events. If the mass distributions of $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) > \frac{\pi}{2}$ events showed a shape too different of that expected for the majority of signal events, then a $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) < \frac{\pi}{2}$ selection might have been necessary. However, as figure 7.6 shows, the shapes of the mass and mass resolution distributions are similar for $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) < \frac{\pi}{2}$ and $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) > \frac{\pi}{2}$ events. The $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) > \frac{\pi}{2}$ distributions are shifted to somewhat smaller values but this is expected since no collinear approximation is possible for them and m_{vis} is used instead. Removing these events is thus unnecessary, and no extra selection is used.

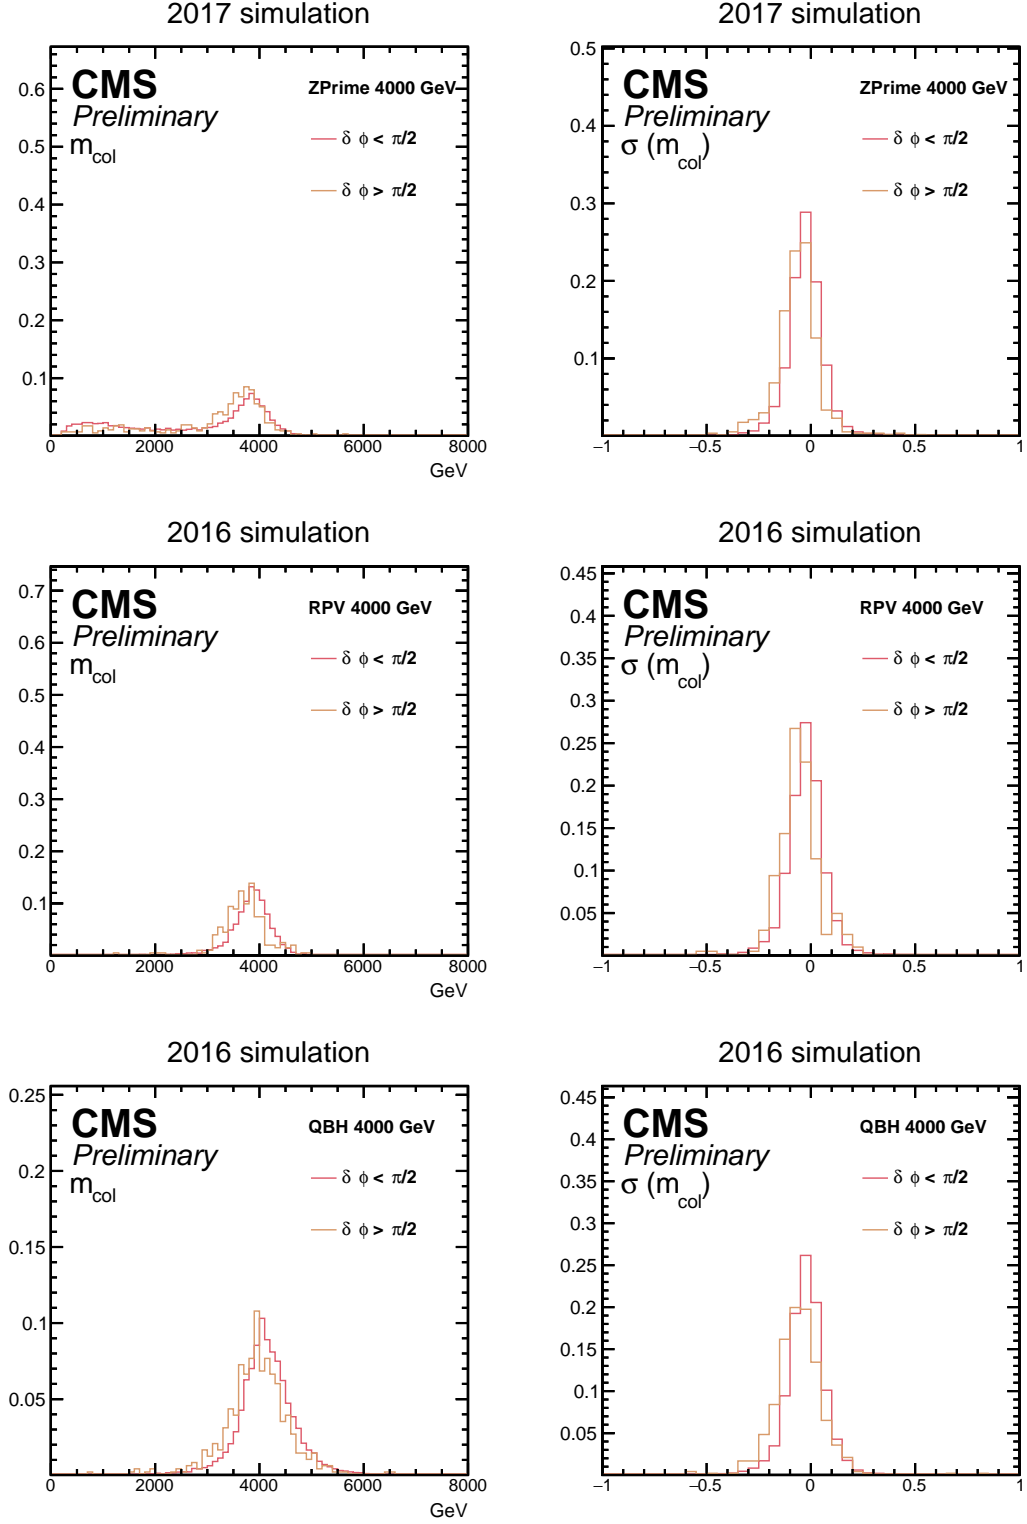


Figure 7.6 – Comparison of $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) > \frac{\pi}{2}$ events to $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) < \frac{\pi}{2}$ events. On the left the reconstructed m_{col} distributions are shown, normalized to 1. On the right, the mass resolution distributions are shown. Note that for $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) > \frac{\pi}{2}$ events the m_{col} definition implies $m_{\text{col}} = m_{\text{vis}}$. The top row shows distributions obtained from a Z' signal sample, the middle row from an RPV SUSY sample, and the bottom row from a QBH sample. The mass point shown for all three signals is 4 TeV, corresponding to the resonance mass for the RPV SUSY $\tilde{\nu}_\tau$ and the Z' boson, and the threshold mass for the QBH.

7.3 Background estimation

The background estimation strategy is similar in the three LFV final states: backgrounds with real leptons are estimated with simulations, while backgrounds with jets faking leptons are estimated with the data-driven fake factor methods detailed in chapter 8. Chapter 6 lists the simulation samples used to simulate real lepton backgrounds, including events in which the lepton flavor is wrongly identified (e.g. a $DY \rightarrow \mu\mu$ event in which a muon is misidentified as an electron).

In the $e\mu$ channel, events with $\text{jet} \rightarrow e$ and $\text{jet} \rightarrow \mu$ fakes are estimated with a data-driven method. In the tau channels, the number of such events is relatively small and they are estimated with simulations, while $\text{jet} \rightarrow \tau$ events exist in considerably larger numbers and these are estimated by a data-driven method.

As explained in chapter 6, simulated events are reweighted in order to correct them for the data-taking conditions and the observed detector response. The $t\bar{t}$ process needs to be particularly well estimated, because it is one of the two most important analysis backgrounds. This is why theory-derived weights are applied on top of the weights which depend on experimental conditions, to bring the shape of the $t\bar{t}$ background in line with NNLO QCD + NLO electroweak corrections [178]. The theory-derived weight w is the geometric average of two weights depending on either the top or the antitop generated transverse momentum, according to the following formula:

$$\begin{aligned} w_{\text{top}} &= e^{(0.0615 - 0.0005 \text{ GeV}^{-1} p_T(\text{top}))} \\ w_{\text{antitop}} &= e^{(0.0615 - 0.0005 \text{ GeV}^{-1} p_T(\text{antitop}))} \\ w &= \sqrt{w_{\text{top}} w_{\text{antitop}}} \end{aligned} \tag{7.4}$$

In order to check how well the $t\bar{t}$ background is estimated, control regions enriched in $t\bar{t}$ events are constructed in the tau channels. This can be done by selecting events with two b-jets, because top quarks almost always decay to a b quark (and a W boson), and there are no other SM backgrounds for this analysis that have such a high proportion of b-jets. The b-jets are required to pass the medium working of the CSV algorithm, CSVv2 for 2016 data and DeepCSV for 2017-18 data. For both the $e\tau$ and $\mu\tau$ channels, two control regions are constructed, one with $m_T < 120 \text{ GeV}$ and another with $m_T > 120 \text{ GeV}$ (technically a subset of the signal region). Figures 7.7 and 7.8 show distributions from the two $t\bar{t}$ control regions in the $\mu\tau$ channel, for the 2018 data-taking period, showing reasonable agreement between observations and expectations. The plots do show a systematic data deficit compared to the simulations, possibly because of a lower b-tagging efficiency in 2018 data which is not accounted for in the simulations. Better agreement is found in other years for both tau channels, as summarized in tables 7.4 and 7.5, validating the $t\bar{t}$ background estimation (see also the plots in appendix B).




Figure 7.7 – $t\bar{t}$ ($m_T > 120$ GeV) control region distributions, $\mu\tau$ channel, 2018. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties, detailed in chapter 9.




Figure 7.8 – $t\bar{t}$ ($m_T < 120$ GeV) control region distributions, $\mu\tau$ channel, 2018. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties, detailed in chapter 9.

Process	2016		2017		2018	
	m_T high	m_T low	m_T high	m_T low	m_T high	m_T low
Observed	636	1227	956	1648	1433	2645
$t\bar{t}$	571	926	824	1333	1450	2308
ST	33	40	45	55	69	96
DY	0	0	2	32	4	48
Diboson	0.4	0.8	0.7	1	1	3
Jet BG	30	260	75	393	99	572
Total BG	634	1218	946	1816	1622	3028

Table 7.4 – Comparison of observed and background event yields in the $e\tau_h t\bar{t}$ control regions.

Process	2016		2017		2018	
	m_T high	m_T low	m_T high	m_T low	m_T high	m_T low
Observed	687	1297	1052	1780	1579	2914
$t\bar{t}$	664	1073	948	1387	1598	2426
ST	34	48	41	55	80	90
DY	0.03	30	2	3	0	34
Diboson	0.3	2	0.4	1	1	3
Jet BG	21	214	81	407	103	621
Total BG	719	1367	1072	1853	1783	3175

Table 7.5 – Comparison of observed and background event yields in $\mu\tau_h t\bar{t}$ control regions.

7.4 Tau identification: comparison between boosted decision trees and DeepTau algorithms

In the beginning of the data analysis work, BDT-based multivariate algorithms were used to identify the hadronic taus. However, as the more sophisticated DeepTau identification algorithms became available, it was important to check if and how much they could improve the analysis in the tau channels. These new algorithms are based on a deep neural network and promise a higher identification efficiency for a similar misidentification rate. The trade-offs between efficiency and misidentification rate can be shown by plotting them as a function of each other in the same two-dimensional graph called a ROC curve.

ROC curves for the BDT and DeepTau algorithms are shown in figure 7.9, for the $e\tau$ and $\mu\tau$ channels. The efficiencies were measured using $Z' \rightarrow e\tau$ and $Z' \rightarrow \mu\tau$ samples, while the misidentification rates were measured in semi-leptonic $t\bar{t}$ simulations in the $e\tau$ channel and W+jets simulations in the $\mu\tau$ channel. The definitions of tau identification efficiency are different in the two channels: the denominator in the $e\tau$ channel simply requires the presence of a hadronic tau matched to a generated lepton, and with $p_T > 30 \text{ GeV}$ and $|\eta| < 2.3$; in the $\mu\tau$ channel the denominator requires the event to pass all analysis selections except for the DeepTau discriminators (and the reconstructed hadronic tau must be matched to a generated hadronic tau). Nevertheless, the efficiency definitions are consistent within each channel, and compared to the old combination of anti-muon, anti-electron, and anti-jet discriminators, the new DeepTau discriminators improve the signal efficiency by 15-20 percentage points for a somewhat smaller jet $\rightarrow \tau_h$ misidentification rate. Switching to the DeepTau ID, and using the tight anti-jet working point, allows for a substantial improvement in the signal sensitivity of

the analysis. Note that there are some differences with the final analysis selections: in figure 7.9 a τ p_T threshold of 30 GeV and the loose anti-muon discriminator were used, rather than the 50 GeV threshold and the tight anti-muon discriminator applied in the final analysis selections. However, these small changes do not invalidate the main conclusion of the study, which is that the DeepTau identification is preferable to the BDT-based algorithms.




Figure 7.9 – Efficiency/misidentification rate trade-offs for the BDT and DeepTau algorithms. The circled gold point corresponds to the old combination of the tight BDT anti-jet (with the MVA 2017v2 training), loose BDT anti-electron discriminators, and the loose cut-based anti-muon selection. The circled light blue point corresponds to the following combination of DeepTau ID algorithms: tight anti-jet, loose anti-electron and loose anti-muon. In the analysis the tight rather than loose anti-muon discriminator is used, but the changes in both signal efficiency and jet misidentification rate are $< 1\%$. In both channels, switching to the DeepTau algorithms allows for both a signal efficiency improvement and a reduction in the misidentification rate.

Chapter summary

The analysis trigger strategy is to select high p_T or high E_T objects, while the offline event selections are kept to the minimum necessary, allowing the search for LFV new physics at the energy frontier to be as model-independent as possible. The specific trigger combinations and event selections vary in the three channels. The tau channels veto events in which the transverse mass of the MET and the light lepton is lower than a 120 GeV threshold, in order to reject backgrounds with $\text{jet} \rightarrow \tau$ fakes. The discriminating variable between signal and backgrounds is chosen to be a mass variable, because the signals are expected to show up in data as localized event excesses in the dilepton mass spectrum. In the $e\mu$ channel this mass variable is simply the dilepton invariant mass, while in the tau channels it is the collinear mass, shown to be more reliable than the alternatives at reconstructing the generated signal event mass. The backgrounds estimated with data-driven fake factor methods are the $\text{jet} \rightarrow e$ and $\text{jet} \rightarrow \mu$ backgrounds in the $e\mu$ channel, and the $\text{jet} \rightarrow \tau$ backgrounds in the tau channels. Backgrounds with real leptons are estimated with simulations, reweighted to correct them for observed detector response and data-taking conditions. Theory-derived weights are also applied to the $t\bar{t}$ background, the description of which is validated in the tau channels by constructing control regions enriched in $t\bar{t}$ events. The DeepTau algorithms are shown to provide an improvement in both the signal efficiency and misidentification rates compared to the old BDT-based tau identification algorithms.

My personal contributions

I was responsible for the full analysis chain in the $\mu\tau$ channel.

- I participated in determining the event selections used to discriminate signal from background in both the tau channels.
- I made the study comparing the different mass variables and showing why the collinear mass is the best choice.
- I made the study comparing $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) > \frac{\pi}{2}$ and $\Delta\phi(\vec{p}_T, \vec{p}_T(\tau_h)) < \frac{\pi}{2}$ events.
- I made the plots in the $t\bar{t}$ control regions for the $\mu\tau$ channel.
- I measured the efficiencies and misidentification rates of the tau identification algorithms in the $\mu\tau$ channel.

Data-driven background estimations

Some background processes are estimated with data-driven methods rather than simulations. These are backgrounds with wrongly-identified leptons, which are particularly hard to simulate. These backgrounds also typically have very large cross sections and would demand an unreasonably large amount of events to be simulated in order to keep the statistical uncertainty low. In the $e\mu$ channel, backgrounds with $\text{jet} \rightarrow e$ and $\text{jet} \rightarrow \mu$ fakes are estimated in a data-driven way, and so are $\text{jet} \rightarrow \tau$ backgrounds in the tau channels. While the estimation methods differ in some details, they follow the same general idea: a fake-lepton-enriched data region is constructed with selections as close as possible to those of the signal region, in this region a fake factor is calculated relating the number of fake lepton candidates passing analysis selections and those which pass looser selections, and finally this fake factor is applied to estimate the number of events with fake leptons in the signal region. Validation tests are then performed to check the reliability of these methods.

This chapter is divided into two sections, with section 8.1 describing the data-driven background estimation performed for the $e\mu$ channel, and section 8.2 explaining in detail the estimation of background events with $\text{jet} \rightarrow \tau$ fakes.

8.1 Data-driven background estimations in the $e\mu$ final state

In the $e\mu$ channel, data-driven methods are used to estimate backgrounds where jets are wrongly identified either as muons or as electrons. Both of these cases correspond to only a few percent of the total background, and thus events in which both the muon and electron are fake are negligible and are not part of the estimation.

8.1.1 Electron and muon fake rates

The shape and yield of the background with $\text{jet} \rightarrow e$ fakes is determined in the following way:

- A data region is constructed with events in which one muon passes all of the muon selections but the electron candidate is required to pass a relaxed set of preselections (shown in table 8.1) but fail the full selections (shown in chapter 5, table 5.1). Requiring the electron candidate to fail the full analysis selections makes this control region independent of the signal region in which the search for new physics is performed.

- The selected region is enriched in $\text{jet} \rightarrow e$ fakes but still contains a significant number of background events with prompt electrons. Simulations of these prompt electron backgrounds are subtracted from the data.
- To get an estimate of $\text{jet} \rightarrow e$ fakes in the signal region, it is necessary to apply the fake factor $f_e = r_e/(1 - r_e)$ to the distributions obtained in the previous step. r_e is the electron fake rate, an energy-dependent function that gives the share of misidentified electron candidates passing the full analysis selections among all fake electrons that pass the loose preselections. The $1/(1 - r_e)$ factor comes from the fact that it is not applied to all events in which electron candidates pass the preselections, but only to those events in which the candidates also fail the full analysis selections.

variable	barrel	endcap
$\sigma_{i\eta i\eta}$	<0.013	<0.034
H/E	<0.15	<0.10
No. of missing hits	<= 1	<= 1
dxy	< 0.02	< 0.05

Table 8.1 – The preselection requirements for the starting point of the fake electron rate calculation.

The fake rate functions (r_e) are taken from the dilepton resonance search at 13 TeV [186, 187, 188]. Table 8.3 shows the function for the HEEPv7.0 ID used in 2016-17 and table 8.4 for the one for the HEEPv7.0-2018Prompt ID used in 2018. As mentioned in chapter 5, the change in the HEEP ID in 2018 was necessary because some modules in one region of the HCAL endcap were non-operational from Run C onwards. This is why the problematic η region of the detector (the “HEM 15-16” region) has a different fake rate function during the data-taking period in which there were issues.

The data-driven estimation of $\text{jet} \rightarrow \mu$ fakes is similar. A region of data is selected in which electrons pass full selections but muon candidates pass looser preselections (shown in table 8.2) and fail to pass the full requirements (described in chapter 5). Simulations of prompt muon backgrounds are subtracted from this data, and the resulting distributions are multiplied by the fake factor $f_\mu = r_\mu/(1 - r_\mu)$ to provide an estimate of events with $\text{jet} \rightarrow \mu$ fakes in the signal region. The fake rate function is also taken from the dilepton resonance search [187, 188].

variable	selection
isGlobalMuon and isTrackerMuon	true
$ d_z $	<1.0
$ d_{xy} $	<0.2
No. of Tracker Layers with Measurements	>5
No. of Valid Pixel Hits	>0
Matching with HLT object (Muon50 TkMu50)	true

Table 8.2 – The preselection requirements for the starting point of the fake muon rate calculation.

8.1.2 Validation with the same sign method

The data-driven estimation of the $\text{jet} \rightarrow e$ and $\text{jet} \rightarrow \mu$ backgrounds by the fake factor method is validated by using an independent procedure to estimate the same backgrounds and

Region	E_T range (GeV)	Functional form
barrel	$35 \leq E_T < 131.6$	$0.14 - 0.0029 \times E_T + 2.56 \cdot 10^{-5} \times E_T^2 - 8.48 \cdot 10^{-8} \times E_T^3$
	$131.6 \leq E_T < 359.3$	$0.002 - 0.00013 \times E_T + 3.5 \cdot 10^{-7} \times E_T^2 - 2.9 \cdot 10^{-10} \times E_T^3$
	$E_T \geq 359.3$	$0.00514 + 4.73 \cdot 10^{-7} \times E_T$
endcap $ \eta < 2.0$	$35 \leq E_T < 125$	$0.1012 - 0.00094 \times E_T + 3.37 \cdot 10^{-6} \times E_T^2$
	$125 \leq E_T < 226.3$	$0.0488 - 11.37 \cdot 10^{-5} \times E_T$
	$E_T \geq 226.3$	$0.0241 - 1.24 \cdot 10^{-6} \times E_T$
endcap $ \eta > 2.0$	$35 \leq E_T < 152$	$0.0622 - 0.00012 \times E_T$
	$E_T \geq 113$	0.0387

Table 8.3 – Functional form of the measured fake rate for HEEPv7.0 (from dielectron resonance search analysis).

Region	E_T range (GeV)	Functional form
barrel before 319077	$35 < E_T < 130$ GeV	$2.02\text{e-}06 \times E_T^2 - 8.43\text{e-}04 \times E_T + 0.087631$
	$130 < E_T < 359.3$ GeV	$-1.34\text{e-}09 \times E_T^3 + 1.11\text{e-}06 \times E_T^2 - 3.11\text{e-}04 \times E_T + 0.0358$
	$E_T > 359.3$ GeV	$2.19\text{e-}07 \times E_T + 0.00599$
endcap $ \eta < 2.0$ before 319077	$35 < E_T < 70$ GeV	$-0.00164 \times E_T + 0.167$
	$70 < E_T < 155$ GeV	$-3.40\text{e-}04 \times E_T + 0.0802$
	$E_T > 155$ GeV	0.0262
endcap $ \eta > 2.0$ before 319077	$35 < E_T < 220$ GeV	$1.34\text{e-}06 \times E_T^2 - 5.84\text{e-}04 \times E_T + 0.106$
	$E_T > 220$ GeV	0.0427
barrel HEM 15-16 region from run 319077	$35 < E_T < 97$ GeV	$-6.41\text{e-}04 \times E_T + 0.0681$
	$E_T > 97$ GeV	0.00846
barrel excluding HEM 15-16 region from run 319077	$35 < E_T < 130$ GeV	$1.18\text{e-}06 \times E_T^2 - 6.89\text{e-}04 \times E_T + 0.0801$
	$130 < E_T < 359.3$ GeV	$-8.21\text{e-}10 \times E_T^3 + 7.09\text{e-}07 \times E_T^2 - 2.09\text{e-}04 \times E_T + 0.0269$
	$E_T > 359.3$ GeV	$7.66\text{e-}06 \times E_T + 0.00291$
endcap $ \eta < 2.0$ HEM 15-16 region from run 319077	$35 < E_T < 100$ GeV	$-4.68\text{e-}04 \times E_T + 0.0663$
	$E_T > 100$ GeV	0.0210277
endcap $ \eta < 2.0$ excluding HEM 15-16 region from run 319077	$35 < E_T < 160$ GeV	$-3.60\text{e-}04 \times E_T + 0.0832$
	$E_T > 160$ GeV	0.0266
endcap $ \eta > 2.0$ HEM 15-16 region from run 319077	$35 < E_T < 1000$ GeV	0.0393
endcap $ \eta > 2.0$ excluding HEM 15-16 region from run 319077	$35 < E_T < 220$ GeV	$1.37\text{e-}06 \times E_T^2 - 5.64\text{e-}04 \times E_T + 0.100$
	$E_T > 220$ GeV	0.0419

Table 8.4 – Functional form of the measured fake rate for HEEPv7.0-2018Prompt (from dielectron resonance search analysis). The “HEM 15-16” region is the η region in which some HCAL endcap modules were non-operational during part of the data-taking (from run 319077, during Run 2018C, onwards).

checking that the two predictions agree. This independent procedure exploits a subset of the signal region in which both leptons are required to have the same sign. Signal events and most Standard Model backgrounds have opposite sign (OS) lepton pairs and thus contribute little to the same sign (SS) control region. However, the charge of a jet misidentified as an electron or muon is approximately uncorrelated to the charge of the real lepton in the event, and thus approximately half of $\text{jet} \rightarrow e$ and $\text{jet} \rightarrow \mu$ background events are in the SS control region. A more complete study, such as that briefly mentioned for the tau identification study in chapter 5, would check the ratio of OS/SS events for the jet background. This ratio is however known to be approximately equal to 1, a value good enough for validation purposes.

There are still some SM background events with prompt leptons in the same sign region, simulations of these events are subtracted from data to get an estimate of the jet background in the SS region. The resulting distributions are doubled to yield the prediction for the jet background in the sign-independent signal region used in the analysis. This prediction can be compared to the one obtained with the fake factor method in order to validate the data-driven background estimation and to provide an estimate for the associated systematic error. Figure 8.1 shows a comparison between the two methods for the three years of LHC Run 2. The same sign method gives results that stay mostly in a 50% range around the results from the fake factor method, and 50% is thus taken as the systematic error on the jet background yield. That is a conservative uncertainty, but since this background is subdominant compared to the $t\bar{t}$ and WW backgrounds, a large uncertainty does not significantly affect the analysis results.

The difference between the fake rates obtained by the two methods could be due to the different flavor compositions of the jets selected by the two methods. As shown in table 8.11, jet flavor has a large influence on the value of the fake rate. Another possible explanation is that the factor of two used to go from the jet background in the SS region to the sign independent jet background is incorrect.

8.2 Backgrounds with jet $\rightarrow \tau$ fakes

For both tau channels ($e\tau$ and $\mu\tau$), the jet faking τ background is estimated from data. Several different approaches were studied in order to check which one produced the best estimation, with the analysis finally settling on the fake factor method described in section 8.2.1.

At first, the possibility of estimating the main jet backgrounds (W+jets and QCD multijets) separately was considered. It would require the use of six data regions (one signal region and five control regions), classified according to the sign of the two leptons and their isolation, as shown in figure 8.2. Here, the tau isolation refers to the tau anti-jet discriminator. Data in which both the light lepton and the tau are anti-isolated contains almost exclusively QCD multijets events, data in which the light lepton is isolated and the tau is anti-isolated contain mostly W+jets and QCD multijets events, while both of these backgrounds are suppressed when both leptons are required to be isolated. The sign selection allows for the creation of a special control region (both leptons isolated, same signs) close to the signal region (both leptons isolated, opposite signs). This control region can then be used to calculate “transfer factors” expressing the changes in the background yields when the lepton isolation requirements are flipped.

This first approach was however abandoned because it required many control regions and was thus too complicated, and in any case it is not necessary to estimate the W+jets and QCD multijets backgrounds separately. The sign selection is also a problem, because although signal events are expected to have opposite sign leptons due to charge conservation, charge misidentifi-

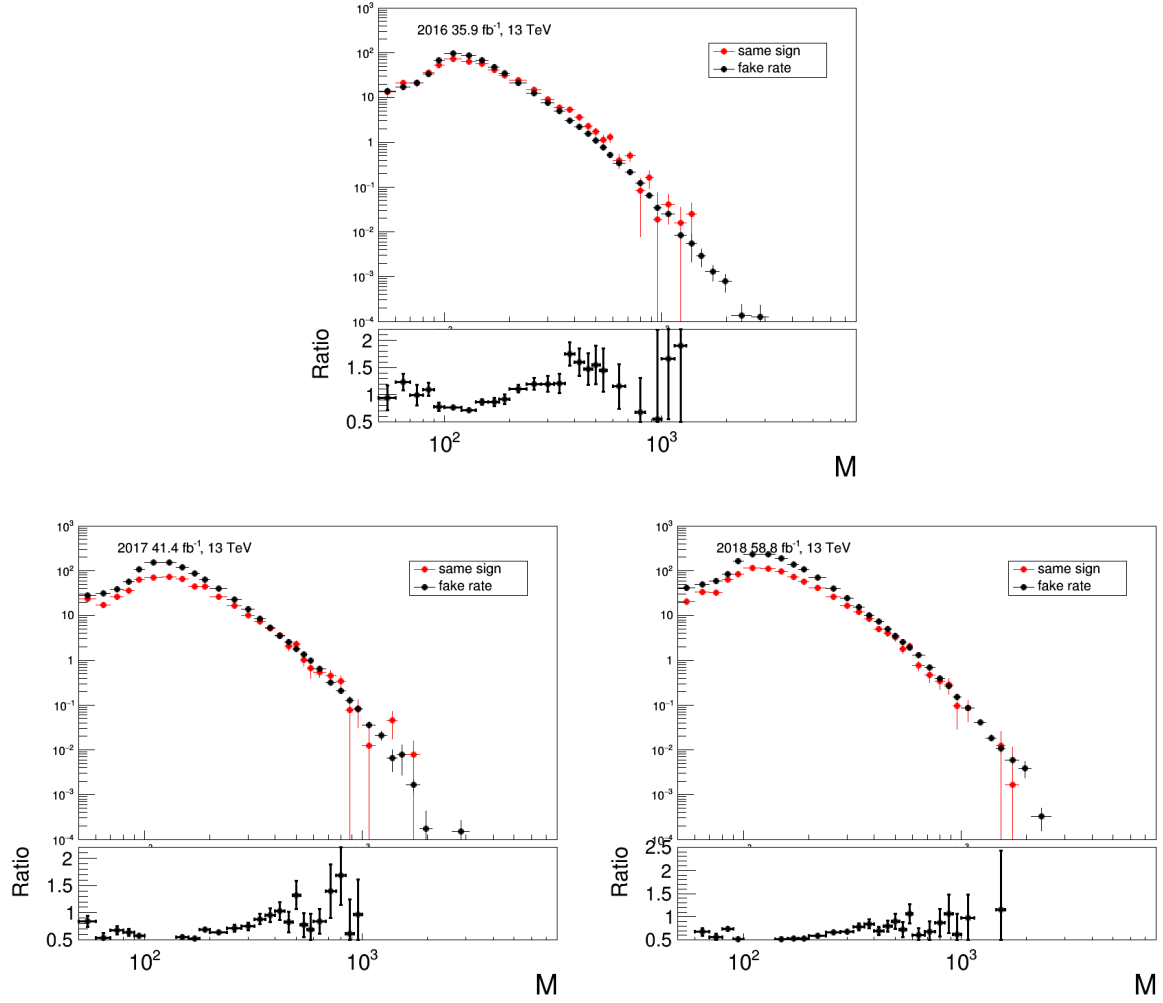


Figure 8.1 – Comparison of the $e\mu$ mass distributions of the jet background (faking electrons or muons) estimated by the fake factor and same sign methods. In all three years, the same sign predictions remain mostly within 50% of the fake rate predictions.

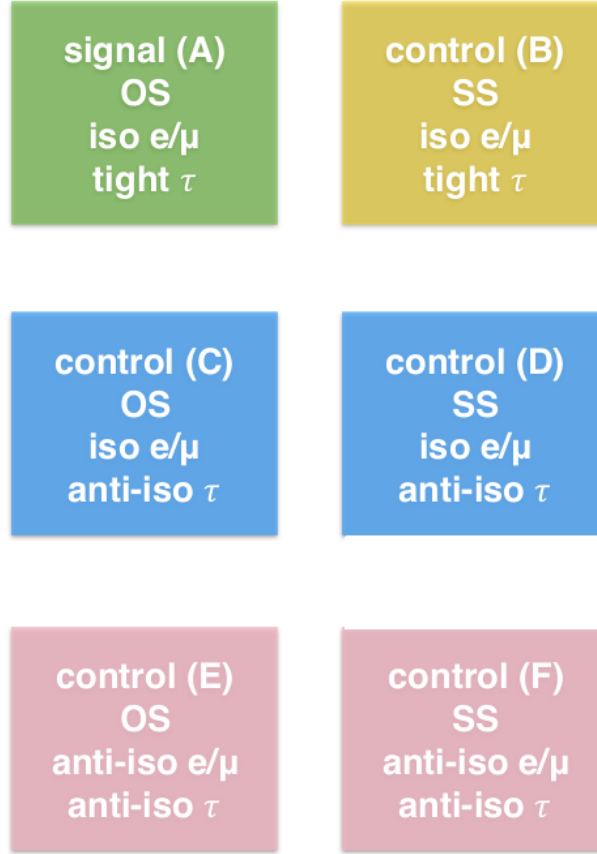


Figure 8.2 – Estimating the QCD multijets and W+jets backgrounds separately, with six control regions characterized by lepton sign and isolation. The QCD multijets background is estimated in the regions where both leptons are anti-isolated and the W+jets background is estimated in the regions where the light lepton is isolated but the tau is anti-isolated. The same sign (SS) regions are used to calculate transfer factors for the yields of the QCD multijets and W+jets backgrounds when lepton isolation is flipped. An F→B transfer factor is calculated for QCD, and a D→B factor for W+jets. These factors are then applied to the distributions obtained in regions E and C to obtain signal region distributions.

cation increases with the particles' transverse energies, and some signal efficiency would be lost. A better way to reject backgrounds with jet $\rightarrow \tau$ fakes while keeping efficiency high is using the $m_T > 120$ GeV selection (cf. chapter 7). W+jets events are expected to have $m_T = 80$ GeV or lower, and there is no particular reason for QCD multijets events to have high m_T , while signal events must have high m_T because m_T is correlated to the invariant mass of the light lepton and the tau. The m_T selection can also be flipped to obtain a control region enriched in backgrounds with jet $\rightarrow \tau$ fakes, with selections that are identical to those of the signal region, except for the m_T requirement. This state of affairs is exploited by the fake factor method.

8.2.1 The fake factor method

The idea of the fake factor method is to consider data events with nothing but jet $\rightarrow \tau$ fakes and then estimate the ratio of such events passing the tight tau anti-jet discriminator to events that fail the tight working point but do pass a loose version of the discriminator. The fake factor is estimated in a control region enriched with jet $\rightarrow \tau$ fakes, ideally a region in

which other processes are negligible. If they are not negligible, then simulations of processes with prompt taus are subtracted from the data. The fake factor may depend on many variables such as the tau p_T and η .

To obtain an estimation of jet $\rightarrow \tau$ fakes in the signal region, the fake factor is then applied to jet $\rightarrow \tau$ fakes failing the tight anti-jet discriminator but passing the loose working point. In equation form, the procedure is:

$$F = \frac{N_{\text{fakes}}^{\text{CR, tight } \tau}}{N_{\text{fakes}}^{\text{CR, anti-tight } \tau}} = \frac{N_{\text{data}}^{\text{CR, tight } \tau} - N_{\text{sim}}^{\text{CR, tight } \tau}}{N_{\text{data}}^{\text{CR, anti-tight } \tau} - N_{\text{sim}}^{\text{CR, anti-tight } \tau}}$$

$$N_{\text{fakes}}^{\text{SR, tight } \tau} = F \cdot N_{\text{fakes}}^{\text{SR, anti-tight } \tau} = F \cdot (N_{\text{data}}^{\text{SR, anti-tight } \tau} - N_{\text{sim}}^{\text{SR, anti-tight } \tau}) \quad (8.1)$$

where F is the fake factor, N_{data} is the number of events in data, N_{sim} is the number of simulated events with prompt taus, and N_{fakes} is the estimated number of events with jet $\rightarrow \tau$ fakes. These event numbers are taken in different data regions, which can be either the background-enriched control region (CR) or the signal region (SR). The CR and SR are further characterized by whether the tau candidate passes the tight anti-jet working point (WP), or fails the tight WP but still passes a loose WP. The latter tau candidates are called “anti-tight”. The goal is to get N_{fakes} in the signal region with the tau passing the tight anti-jet discriminator.

This procedure can be repeated in many control regions, and observed variations in the calculated fake factors can give some insight into the systematic errors associated to this data-driven background estimation.

8.2.2 Fake factors from different control regions

Fake factors were calculated in a few different control regions:

- The low m_T region, which uses the same selections as the signal region, except for the m_T requirement, which is flipped to $m_T < 120$ GeV. This control region can be further split into a region where the tau and the light lepton (e or μ) have the same signs (SS) and another where they have opposite signs (OS).
- The DY+jets control region, which is constructed as follows:
 - Trigger: HLT_IsoMu24 OR HLT_IsoTkMu24 (2016 data only).
 - Events should have exactly two muons which must pass medium ID, $p_T > 26$ GeV, $|\eta| < 2.4$, PF isolation (delta-beta corrected) < 0.15 , $|d_{xy}| < 0.045$, $|d_z| < 0.2$,
 - $\Delta R(\mu, \mu) > 0.5$,
 - $70 < M_{\mu\mu} < 110$ GeV
 - Presence of a tau candidate, with $\Delta R(\tau, \mu) > 0.5$, for both muons.

For each of these control regions, fake factors are calculated as shown in equation 8.1, but note that the DY+jets region is so thoroughly dominated by fakes that it is unnecessary to subtract any simulation of prompt taus from the data. At first, these fake factors are characterized by tau candidate p_T and whether the tau candidate is located in the barrel or one of the endcaps. Plots comparing the fake factors obtained in the three different regions are shown in figure 8.3, with 2016 data. In both the barrel and endcaps, the fake factors in the low m_T regions are up to 40% different compared to those calculated in the DY+jets region.

Figures 8.4-8.5 show fake factors for the different control regions when they are also characterized by the decay mode of the tau candidate. The statistical uncertainties are higher because

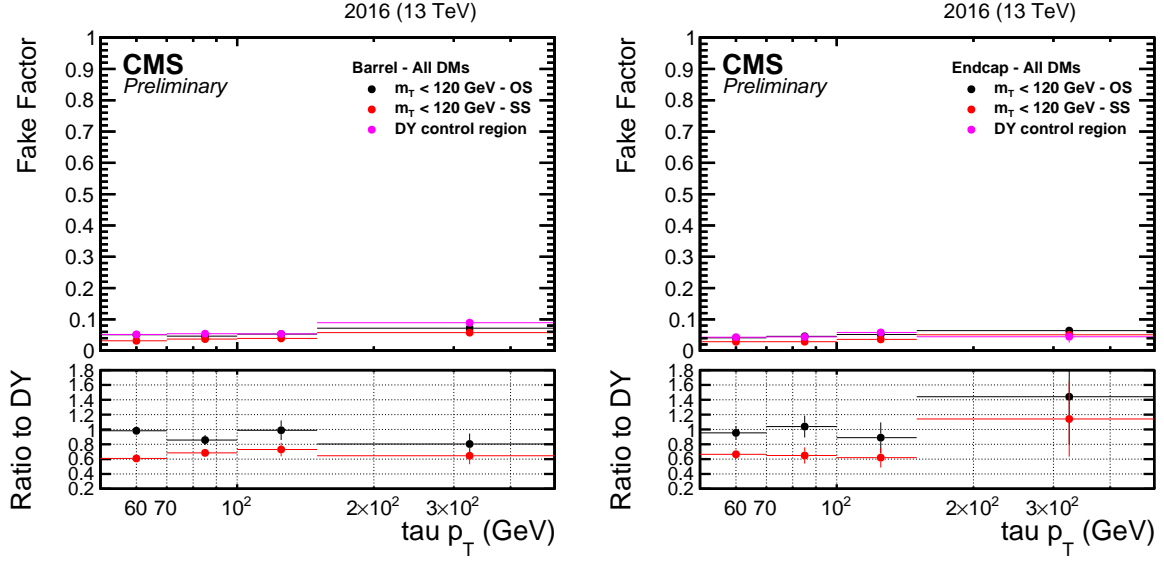


Figure 8.3 – jet $\rightarrow \tau$ fake factors as a function of the p_T of the tau candidate, for three control regions: low m_T OS, low m_T SS and DY+jets. On the left, the tau candidate is in the barrel, while on the right the tau candidate is in one of the endcaps. Error bars are statistical only.

there are fewer events in each bin. Fake factors can vary substantially from one decay mode to the other. The lowest fake factors are found for “3 prong + 1 π^0 ” taus (DM11).

The low m_T control region is the one that is closest to the selections from the signal region, and it also benefits from lower statistical uncertainty than the DY+jets region. And since the signal region has no sign selection, the fake factors used in the analysis are calculated in the low m_T region with no sign selection, rather than taking only OS or SS events. The fake factor dependence on the control region in which it is calculated provides an estimate of the systematic uncertainty associated with the jet $\rightarrow \tau$ fakes background, and so does the study of the dependence of the fake factor on the tau jet flavor (cf. section 8.2.5).

8.2.3 Fake factor dependence on the light lepton p_T

Applying the fake factors calculated in the previous section, which depend only on the tau candidate p_T and its position on the barrel/endcap, yields a strong dependence on the light lepton (e/μ) p_T for the ratio of observed to expected events. This suggests that the fake factor depends on an extra variable, and not including that variable is skewing the distributions of the tau fakes background.

It turns out that this dependence of the fake factor on the light lepton p_T can be accounted for by the fake factor dependence on the ratio of the tau candidate p_T to the p_T of its parent jet. The tau candidates do not necessarily match perfectly the AK4 PF jets they are reconstructed from, sometimes only a subset of the parent jet is used to reconstruct the tau candidate. When the ratio of the tau p_T over the jet p_T is low, the tau candidate is less likely to pass tight isolation requirements, as there is a lot of activity in the detector near the tau, coming from the parts of the jet which were not used to construct the tau. Thus, for a tau candidate with a given p_T , its fake factor will tend to be lower if its parent jet has a much higher p_T , i.e. if the tau over jet p_T ratio is low. The light lepton p_T has a stronger correlation with the parent jet p_T than with the p_T of the tau candidate itself, so the dependence of the fake factor on the ratio implies a dependence on the light lepton p_T .

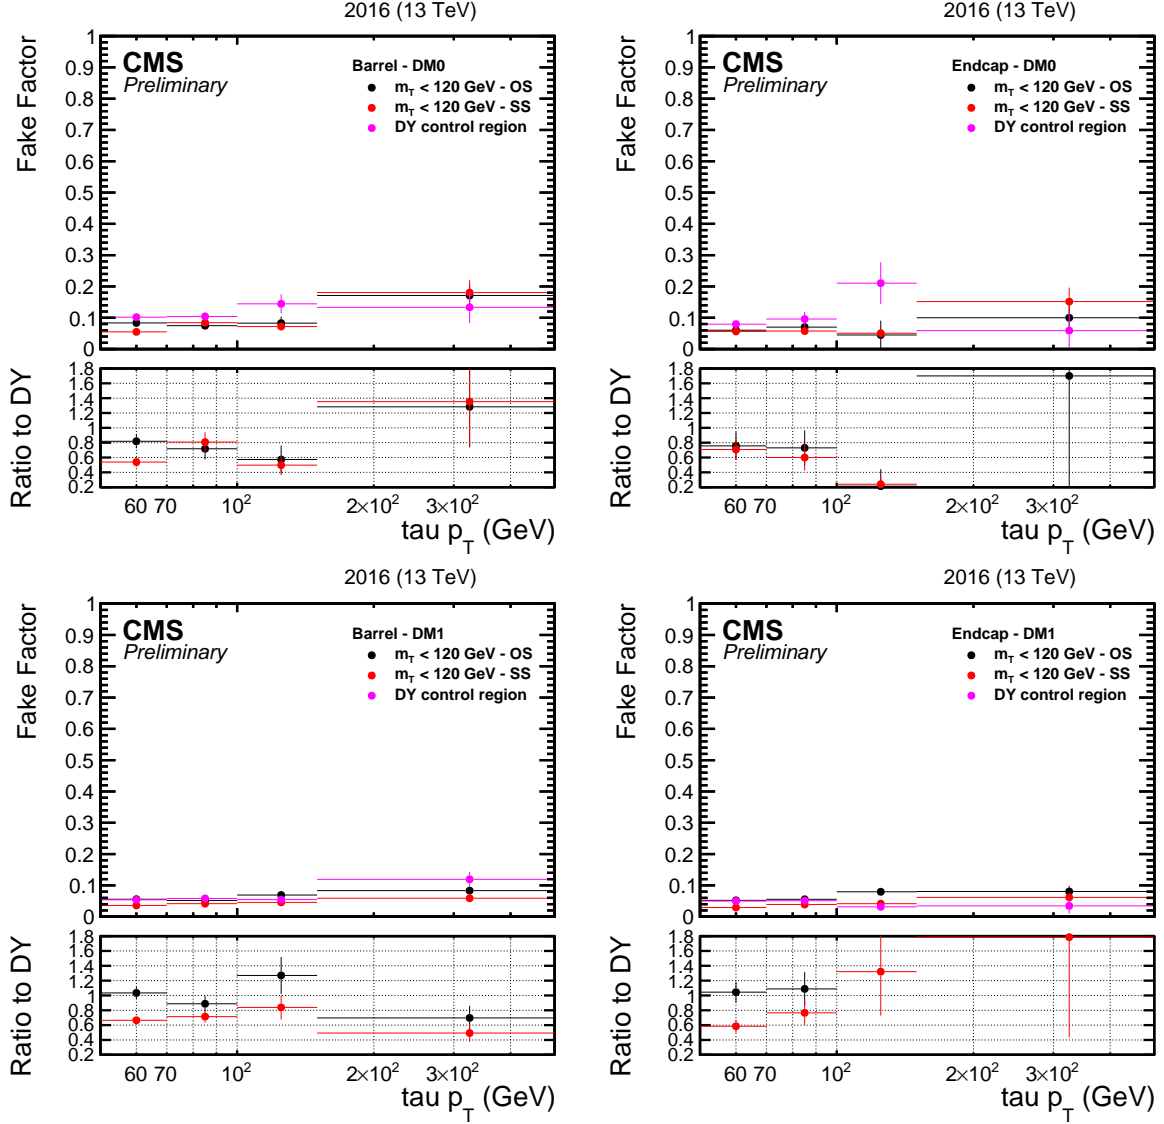


Figure 8.4 – jet $\rightarrow \tau$ fake factors as a function of the p_T of the tau candidate, for three control regions: low m_T OS, low m_T SS and DY+jets. The fake factors are characterized by both the tau candidate decay mode (DM) and its position in the barrel or one of the endcaps. Plots in the first row show tau candidates reconstructed with DM0 (1 prong + 0 π^0), and the second row shows DM1 (1 prong + 1 π^0).

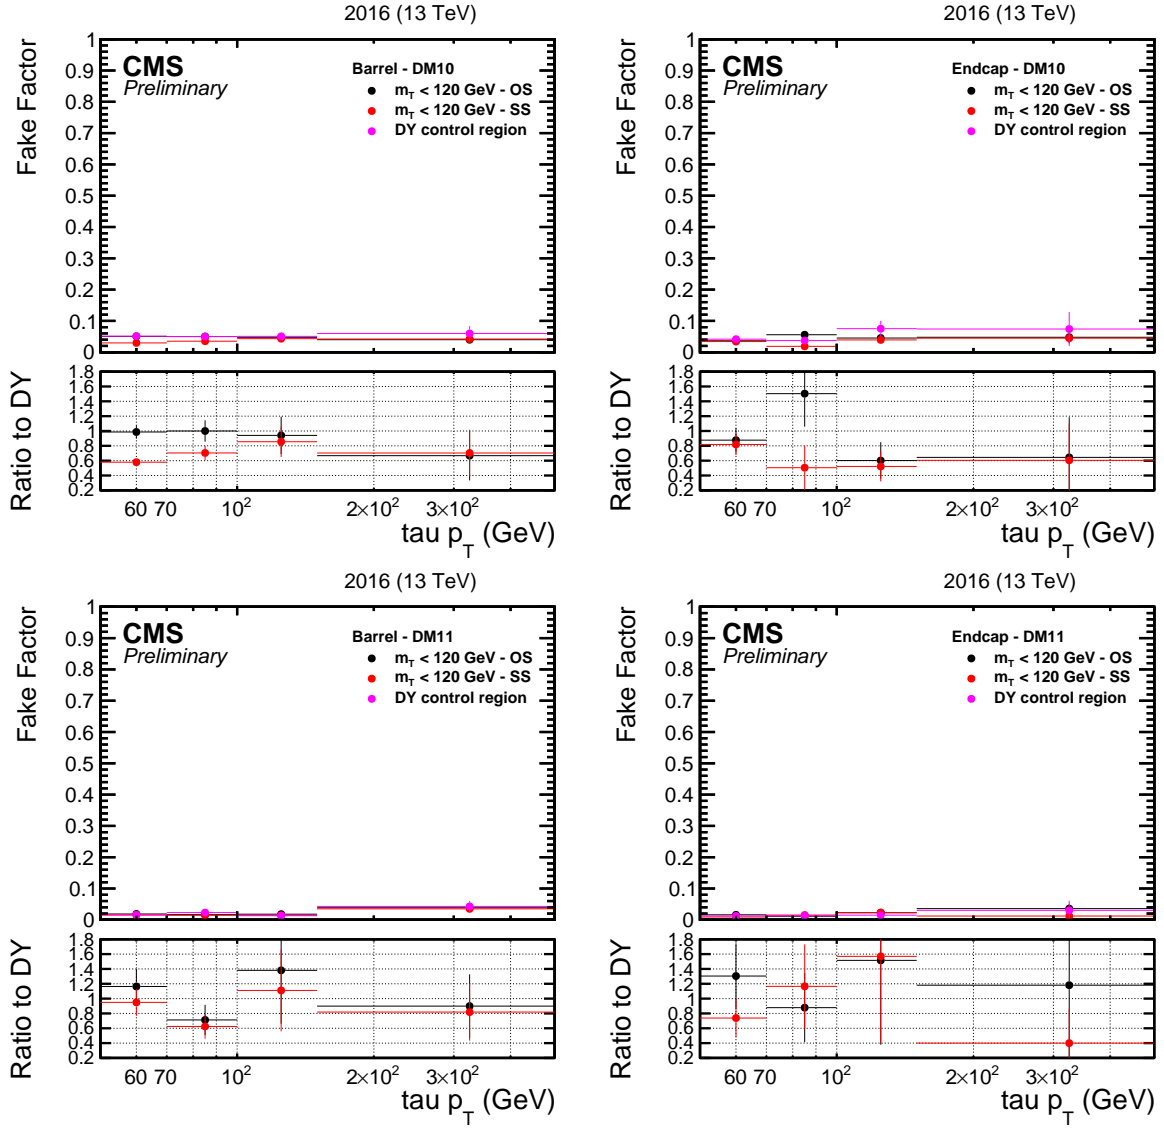


Figure 8.5 – Same as figure 8.4, but with different decay modes shown. Plots in the first row show tau candidates reconstructed with DM10 (3 prong + 0 π^0), the second row shows DM11 (3 prong + 1 π^0).

In the most general case, one could estimate fake factors fully differential in tau candidate p_T and in the tau candidate over jet p_T ratio (called simply 'ratio' from now on). But assuming that the fake factors for different tau p_T 's are affected in the same way by the same ratios, the fake factor could be *factorized* into two terms, $F \cdot R$:

- $F(p_T)$ depends only on the tau candidate's p_T .
- $R(r)$ is a correction factor, averaging 1, which depends only on the ratio r .

This would reduce considerably the number of bins the data is sliced into, and thus reduce the statistical uncertainty on the fake factors and correction factors obtained in each bin. A study was done in the $\mu\tau$ final state using the 2016v1 BDT-based tau ID to check which fake factor estimation method is more effective. Figure 8.6 shows the different correction factors obtained for different tau p_T ranges, and illustrates the limits of this factorization assumption. For ratios close to 1, they are off by as little as 10%, but for ratios close to 0.5, there is a factor of two between the correction factors for the lowest and highest tau p_T ranges. A clear trend in tau p_T can be seen: the correction factors are systematically lower as the tau p_T increases. Due to these large discrepancies in the correction factors, we cannot justify using p_T -independent ratio correction factors. Fake factors fully differential in tau- η , tau- p_T and the ratio are necessary.

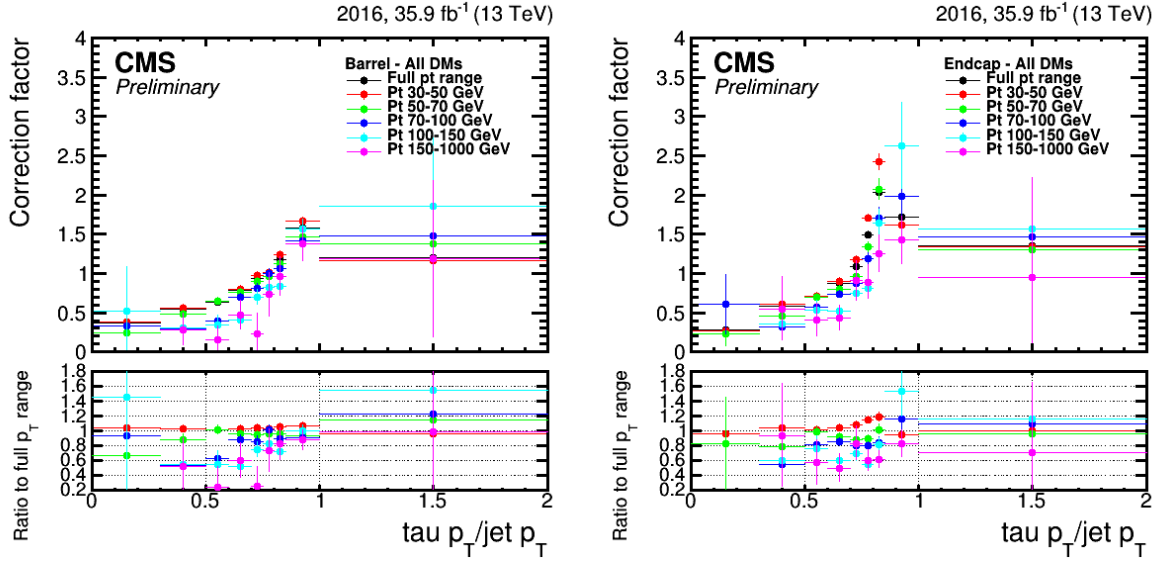


Figure 8.6 – Correction factors as a function of the p_T ratio of the tau candidate over its parent jet, for different tau p_T ranges. On the left, the tau candidate is required to be in the barrel, on the right it's required to be in one of the endcaps. Error bars are statistical only. This study was done with the 2016v1 BDT-based tau ID.

8.2.4 Validation using simulated events

The fake factors used for the jet $\rightarrow \tau$ background estimation in the $e\tau$ and $\mu\tau$ final states are given in tables 8.5-8.10, for each of the three analysis years. The numbers obtained for each bin do not vary much from channel to channel and year to year.

$\tau_h p_T$	τ_h over jet p_T ratio	fake factor (barrel)	fake factor (endcap)
50-80	0.0-0.5	0.005 ± 0.001	0.005 ± 0.002
	0.5-0.6	0.011 ± 0.001	0.011 ± 0.001
	0.6-0.65	0.021 ± 0.002	0.013 ± 0.002
	0.65-0.7	0.024 ± 0.002	0.025 ± 0.002
	0.7-0.75	0.032 ± 0.002	0.039 ± 0.002
	0.75-1.0	0.069 ± 0.002	0.066 ± 0.004
	1.0-3.0	0.037 ± 0.006	0.042 ± 0.008
80-150	0.0-0.5	0.004 ± 0.002	0.005 ± 0.003
	0.5-0.6	0.009 ± 0.002	0.009 ± 0.004
	0.6-0.65	0.014 ± 0.003	0.01 ± 0.004
	0.65-0.7	0.018 ± 0.002	0.023 ± 0.004
	0.7-0.75	0.026 ± 0.003	0.033 ± 0.004
	0.75-1.0	0.059 ± 0.003	0.062 ± 0.005
	1.0-3.0	0.053 ± 0.014	0.078 ± 0.02
150-1000	0.0-0.7	0.017 ± 0.005	0.003 ± 0.009
	0.7-1.0	0.069 ± 0.006	0.069 ± 0.008
	1.0-3.0	0.046 ± 0.032	0.0 ± 0.063

Table 8.5 – **2016** Fake factors in bins of $\tau_h p_T$ and $\tau_h p_T$ over jet p_T ratio for the $\mu\tau_h$ final state.

$\tau_h p_T$	τ_h over jet p_T ratio	fake factor (barrel)	fake factor (endcap)
50-80	0-0.5	0.0054 ± 0.001	0.0043 ± 0.0015
	0.5-0.6	0.01 ± 0.001	0.012 ± 0.0014
	0.6-0.65	0.018 ± 0.0013	0.018 ± 0.0016
	0.65-0.7	0.023 ± 0.0012	0.022 ± 0.0016
	0.7-0.75	0.03 ± 0.0016	0.039 ± 0.0022
	0.75-1	0.073 ± 0.0021	0.074 ± 0.0038
	1-5	0.046 ± 0.0042	0.035 ± 0.0063
80-150	0-0.5	0.0067 ± 0.0024	0.0092 ± 0.0043
	0.5-0.6	0.0093 ± 0.0025	0.011 ± 0.0036
	0.6-0.65	0.021 ± 0.003	0.018 ± 0.0089
	0.65-0.7	0.025 ± 0.0029	0.028 ± 0.0037
	0.7-0.75	0.028 ± 0.0036	0.038 ± 0.0037
	0.75-1	0.065 ± 0.0033	0.061 ± 0.0047
	1-5	0.06 ± 0.011	0.052 ± 0.013
150-1000	0-0.7	0.016 ± 0.005	0.033 ± 0.012
	0.7-1	0.087 ± 0.0074	0.069 ± 0.0098
	1-5	0.0038 ± 0.038	0.037 ± 0.043

Table 8.6 – **2016** Fake factors in bins of $\tau_h p_T$ and $\tau_h p_T$ over jet p_T ratio for the $e\tau_h$ final state.

$\tau_h p_T$	τ_h over jet p_T ratio	fake factor (barrel)	fake factor (endcap)
50-80	0.0-0.5	0.005 ± 0.001	0.008 ± 0.001
	0.5-0.6	0.011 ± 0.001	0.018 ± 0.001
	0.6-0.65	0.022 ± 0.001	0.024 ± 0.002
	0.65-0.7	0.03 ± 0.001	0.036 ± 0.002
	0.7-0.75	0.044 ± 0.001	0.063 ± 0.003
	0.75-1.0	0.087 ± 0.002	0.089 ± 0.006
	1.0-3.0	0.056 ± 0.005	0.038 ± 0.006
80-150	0.0-0.5	0.009 ± 0.002	0.009 ± 0.004
	0.5-0.6	0.011 ± 0.002	0.021 ± 0.004
	0.6-0.65	0.021 ± 0.003	0.022 ± 0.004
	0.65-0.7	0.029 ± 0.002	0.029 ± 0.004
	0.7-0.75	0.04 ± 0.003	0.05 ± 0.004
	0.75-1.0	0.079 ± 0.003	0.068 ± 0.007
	1.0-3.0	0.072 ± 0.013	0.04 ± 0.014
150-1000	0.0-0.7	0.027 ± 0.005	0.02 ± 0.007
	0.7-1.0	0.076 ± 0.006	0.075 ± 0.01
	1.0-3.0	0.022 ± 0.03	0.048 ± 0.03

Table 8.7 – **2017** Fake factors in bins of $\tau_h p_T$ and $\tau_h p_T$ over jet p_T ratio for the $\mu\tau_h$ final state.

$\tau_h p_T$	τ_h over jet p_T ratio	fake factor (barrel)	fake factor (endcap)
50-80	0-0.5	0.0042 ± 0.00086	0.0088 ± 0.0017
	0.5-0.6	0.015 ± 0.00094	0.018 ± 0.0014
	0.6-0.65	0.022 ± 0.0011	0.029 ± 0.0016
	0.65-0.7	0.032 ± 0.0012	0.04 ± 0.0018
	0.7-0.75	0.046 ± 0.0013	0.068 ± 0.0025
	0.75-1	0.097 ± 0.0017	0.089 ± 0.005
	1-5	0.054 ± 0.0061	0.047 ± 0.0071
80-150	0-0.5	0.0032 ± 0.0019	0.0068 ± 0.0037
	0.5-0.6	0.012 ± 0.0023	0.018 ± 0.0032
	0.6-0.65	0.021 ± 0.0028	0.018 ± 0.0035
	0.65-0.7	0.03 ± 0.0024	0.035 ± 0.0035
	0.7-0.75	0.037 ± 0.0025	0.053 ± 0.0038
	0.75-1	0.082 ± 0.0027	0.062 ± 0.006
	1-5	0.07 ± 0.016	0.025 ± 0.017
150-1000	0-0.7	0.017 ± 0.0044	0.045 ± 0.011
	0.7-1	0.087 ± 0.0066	0.073 ± 0.012
	1-5	0.12 ± 0.049	0.087 ± 0.048

Table 8.8 – **2017** Fake factors in bins of $\tau_h p_T$ and $\tau_h p_T$ over jet p_T ratio for the $e\tau_h$ final state.

$\tau_h p_T$	τ_h over jet p_T ratio	fake factor (barrel)	fake factor (endcap)
50-80	0.0-0.5	0.007 ± 0.001	0.005 ± 0.001
	0.5-0.6	0.013 ± 0.001	0.017 ± 0.001
	0.6-0.65	0.018 ± 0.001	0.022 ± 0.002
	0.65-0.7	0.027 ± 0.001	0.031 ± 0.002
	0.7-0.75	0.04 ± 0.001	0.057 ± 0.002
	0.75-1.0	0.08 ± 0.002	0.088 ± 0.005
	1.0-3.0	0.057 ± 0.004	0.046 ± 0.005
80-150	0.0-0.5	0.007 ± 0.002	0.005 ± 0.003
	0.5-0.6	0.012 ± 0.002	0.013 ± 0.003
	0.6-0.65	0.017 ± 0.002	0.017 ± 0.003
	0.65-0.7	0.024 ± 0.002	0.026 ± 0.003
	0.7-0.75	0.039 ± 0.002	0.052 ± 0.004
	0.75-1.0	0.064 ± 0.002	0.078 ± 0.009
	1.0-3.0	0.067 ± 0.01	0.042 ± 0.01
150-1000	0.0-0.7	0.019 ± 0.004	0.037 ± 0.009
	0.7-1.0	0.084 ± 0.005	0.074 ± 0.009
	1.0-3.0	0.04 ± 0.025	0.041 ± 0.038

Table 8.9 – **2018** Fake factors in bins of $\tau_h p_T$ and $\tau_h p_T$ over jet p_T ratio for the $\mu\tau_h$ final state.

$\tau_h p_T$	τ_h over jet p_T ratio	fake factor (barrel)	fake factor (endcap)
50-80	0-0.5	0.0053 ± 0.00084	0.0063 ± 0.0013
	0.5-0.6	0.012 ± 0.00073	0.016 ± 0.0011
	0.6-0.65	0.018 ± 0.00092	0.024 ± 0.0013
	0.65-0.7	0.029 ± 0.00092	0.038 ± 0.0015
	0.7-0.75	0.04 ± 0.001	0.059 ± 0.0019
	0.75-1	0.084 ± 0.0014	0.094 ± 0.0037
	1-5	0.053 ± 0.0042	0.047 ± 0.0052
80-150	0-0.5	0.0064 ± 0.0018	0.013 ± 0.004
	0.5-0.6	0.015 ± 0.002	0.016 ± 0.0028
	0.6-0.65	0.02 ± 0.0022	0.018 ± 0.0033
	0.65-0.7	0.031 ± 0.0021	0.031 ± 0.003
	0.7-0.75	0.038 ± 0.0021	0.045 ± 0.003
	0.75-1	0.073 ± 0.0022	0.068 ± 0.0047
	1-5	0.085 ± 0.011	0.061 ± 0.012
150-1000	0-0.7	0.02 ± 0.0049	0.03 ± 0.0071
	0.7-1	0.082 ± 0.0054	0.062 ± 0.011
	1-5	0.04 ± 0.048	0.14 ± 0.048

Table 8.10 – **2018** Fake factors in bins of $\tau_h p_T$ and $\tau_h p_T$ over jet p_T ratio for the $e\tau_h$ final state.

The fake factor numbers are validated by applying them to W+jets simulations, and then two distributions are compared:

- W+jets events with tau candidates passing the tight anti-jet selection.
- W+jets events with anti-tight tau candidates (failing the tight selection but passing a looser one). These events are then rescaled by the fake factors.

If fake factors do what they are supposed to do, the two distributions should be compatible. Figure 8.7 shows such a comparison using the 2017 fake factors on 2017 W+jets simulations, for the tau p_T and collinear mass variables, for both of the tau channels. The distributions are indeed compatible, given statistical uncertainties.

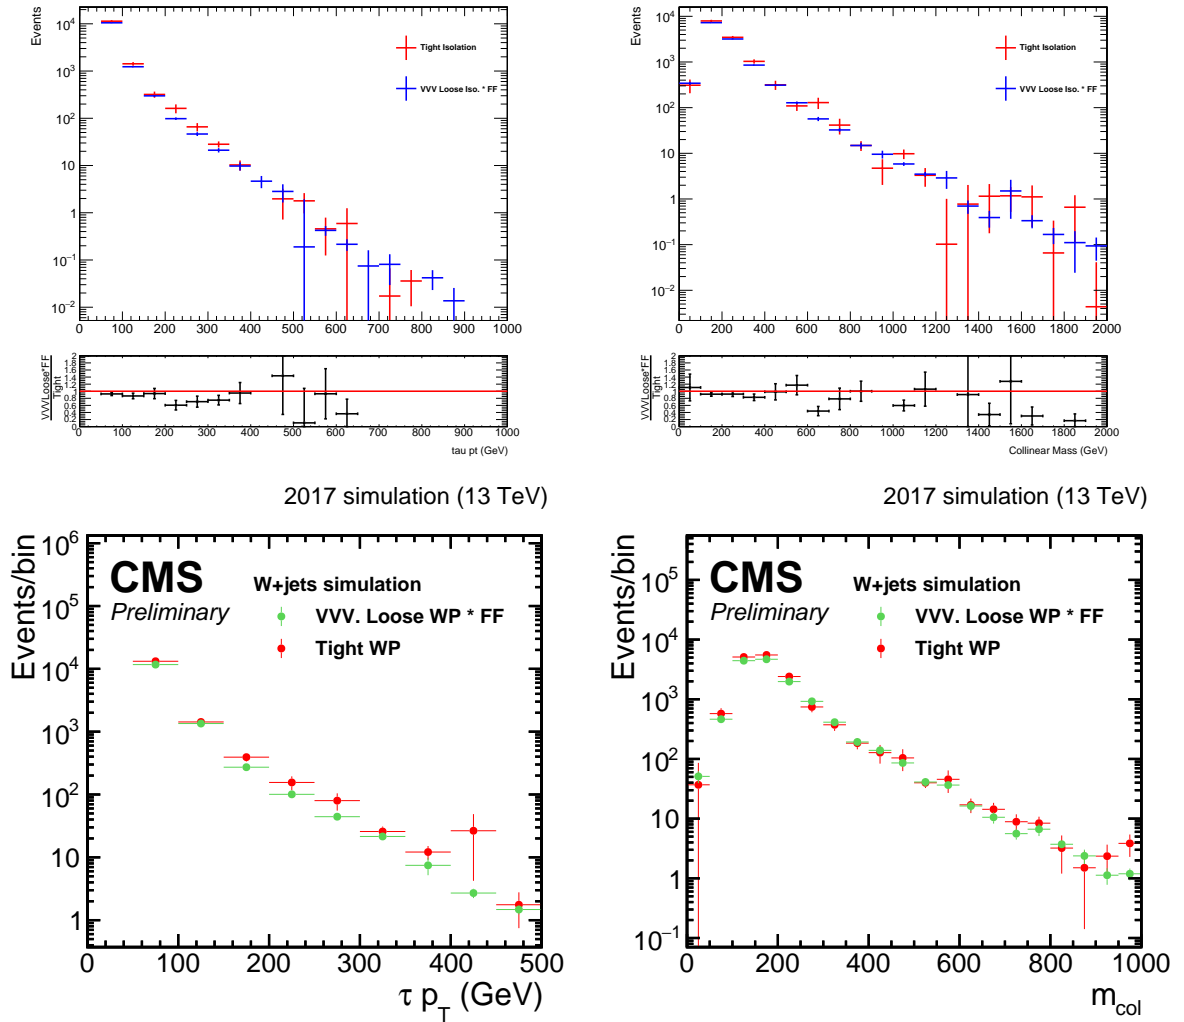


Figure 8.7 – Validation of the 2017 fake factors using 2017 W+jets simulations. On the top row, the $e\tau$ channel is shown, while the $\mu\tau$ channel is in the bottom row. Plots on the left show the tau candidate p_T distributions, and the collinear mass distributions are on the right. Uncertainties are statistical only.

8.2.5 Fake factor dependence on jet flavor

In an effort to understand the systematic biases that might affect the fake factor calculations, we studied the dependence of the fake factor on the flavor of the jet from which fake tau candidates are reconstructed. Jet flavor is simply the flavor of the parton (quark or gluon) which initiated the jet. Since top quarks decay too fast to form proper jets, only six flavors are possible: gluons, or the five lightest quarks (up, down, strange, charm, bottom). Simulations rather than data were used for this study, since jet flavor identifiers only exist for some of the heavy flavors.

Fake factors for each of the six flavors were calculated in 2017 W+jets simulations, and they are compared in table 8.11. The fake factors vary by a lot, the value for the up quark is almost $4\times$ that of the bottom quark. The variations of up to 40 % observed in the fake factors calculated in different control regions (cf. section 8.2.2) could be due to the different flavor composition of these regions.

An uncertainty of 50 %, if applied to the fake factors of the two most common flavors, the up and down quarks, would cover the fake factors of all other flavors, except for the bottom quark. This would be a conservative uncertainty since the extreme case where all jets have the same flavor is unlikely. A less conservative approach would be to consider different realistic flavor compositions for the jets and calculate the average fake rate in each case. However, as the yield of the $\text{jet} \rightarrow \tau$ background is low for high masses, the impact of this uncertainty on the final results is small, and the conservative uncertainty of 50 % is chosen for this background.

Jet flavor	Fake factor
down	0.062 ± 0.006
up	0.077 ± 0.005
strange	0.039 ± 0.01
charm	0.068 ± 0.007
bottom	0.020 ± 0.007
gluon	0.044 ± 0.011

Table 8.11 – Fake factor dependence on jet flavor, calculated in 2017 W+jets samples, for tau candidates with $p_T \in [50, 500]$ GeV. Uncertainties are statistical only.

Chapter summary

Some backgrounds with wrongly-identified leptons are estimated with data-driven methods because the simulations for these processes are unreliable. “Fake factor” methods estimate the $\text{jet} \rightarrow \tau$ background in the tau channels, and the $\text{jet} \rightarrow e$ and $\text{jet} \rightarrow \mu$ backgrounds in the $e\mu$ channel. The $e\mu$ channel uses fake factors already calculated for the dilepton same-flavor resonance search. This estimation of fake leptons is validated by an alternative method using a subset of the signal region with same sign events. The tau channels use novel fake factors calculated specifically for this LFV analysis. The fake factor dependence on the region in which they are calculated, on the light lepton p_T and on the tau candidate jet flavor have all been thoroughly studied, allowing for more accurate estimates of the $\text{jet} \rightarrow \tau$ background and an understanding of the underlying systematic uncertainties. The $\text{jet} \rightarrow \tau$ fake factors were validated by applying them to simulations of W +jets events.

My personal contributions

I had many contributions to the data-driven background estimation in the tau channels:

- I performed the studies checking the best way to calculate the $\text{jet} \rightarrow \tau$ fake factors.
- I compared various ways of taking into account the tau over parent jet p_T ratio.
- I estimated the systematic uncertainty by considering different control regions and their dependence on jet flavor.
- I computed the fake factors in the $\mu\tau$ channel and I validated their performance on simulations.

Signal region events and statistical analysis

The previous chapters have described what data is used to search for new LFV physics, what subset of the data is selected, how signal events are simulated, and how Standard Model backgrounds are estimated, whether by simulation or data-driven methods. In this chapter, data distributions are compared to background expectations to check how well they agree. In particular, we look for any localized excess of data events in the mass distribution (collinear mass in the tau channels) which could signal the existence of new physics. For these comparisons to be meaningful, an understanding of all systematic uncertainties affecting the three LFV channels of the analysis is necessary. If there is no data excess relative to background expectations, above and beyond what is expected from the systematic uncertainties, a statistical analysis is performed to determine limits on the existence of new physics at the 95% confidence level.

Section 9.1 reviews all systematic uncertainties affecting the analysis. Section 9.2 shows comparisons between data and background distributions of various lepton kinematical variables and the mass variables which are the discriminating variables of the analysis. Since no data excess relative to background expectation is found, section 9.3 explains the statistical analysis that determines 95% confidence level limits on three specific new LFV physics models: Z' boson, quantum black holes, and RPV SUSY. The chapter closes with a discussion of the results and prospects for future analyses in section 9.4.

9.1 Systematic uncertainties

The analysis backgrounds are only estimated with finite precision, and so are the signal events we are looking for. The first kind of uncertainty is statistical: for simulation samples the finite number of generated events puts a limit on how precisely the corresponding process is estimated. Statistical uncertainties are well understood and straightforward to calculate, but there are other sources of error, which are harder to estimate, called systematic uncertainties. These uncertainties apply to nuisance parameters, i.e. experimental and theoretical quantities that are only known with finite precision and which are not the main parameter of interest of the analysis (the cross section of new physics processes). The nuisance parameters correspond to biases in the experimental apparatus (detector response, calibration) or in theory-derived values such as the cross sections of individual processes. Systematic uncertainties can affect the shape of the distributions, their normalization, or both. Normalization uncertainties are usually described by log-normal uncertainties, while for shape uncertainties two distributions showing $\pm 1 \sigma$ variation with respect to the nominal distributions are provided.

Log-normal distributions are used for nuisance parameters that can only take positive values. Normal distributions would assign non-zero probabilities to negative values of the parameter, and trying to truncate the Gaussian so that it only covers positive values can lead to inconsistent results [189]. In a log-normal distribution, it is the logarithm of the parameter which is assumed to follow a normal distribution.

The full list of systematic uncertainties used in the analysis is given below:

- A minimum bias cross section¹ of 69.2 mb is used for pileup reweighting. An uncertainty of ± 5 % on the minimum bias cross section is used, a value recommended by the LUMI POG. This uncertainty is incorporated as a shape uncertainty and is considered for all signals and backgrounds.
- The uncertainty on integrated luminosity is considered to be ± 2.5 %, a value which is also recommended by the LUMI POG. It is incorporated as a log-normal normalization uncertainty and is considered for all simulated signals and backgrounds, though not for data-driven backgrounds.
- The uncertainty associated with the choice of PDF sets for signal samples is taken into account as a shape uncertainty. For the Z' signal, the NNPDF3.1 weights are used², with 101 MC replicas to estimate the uncertainties. Only one MC replica is available for each sample, but by using the LHAPDF6 package [190] it is possible to reweight the events to obtain the extra MC replicas. For each mass point, the MC replicas are classified according to their cross sections, and then the MC replicas with the 16th and 84th highest cross sections are taken as the “up” and “down” variations of the PDF uncertainty. For high mass samples, some of the MC replicas have negative cross sections, those are removed from the pool before classifying the MC replicas by cross section. For the RPV and QBH signals, the reweighting procedure is unstable, so the uncertainties calculated in the 2016 $e\mu$ analysis [82] are used here too, by fitting a fourth degree polynomial depending on the LHE-level dilepton mass to the shape of the uncertainty reported in the analysis note. The polynomial evaluated at the event’s dilepton mass gives the PDF uncertainty for that event.
- Muon energy scale: the systematic uncertainty on the muon p_T scale is taken into account as a shape uncertainty. A curvature bias as a function of muon η and ϕ is measured in collision data by the Muon POG, with the generalized endpoint method [191]. The muon momentum scale is then shifted by the curvature bias per TeV. This uncertainty is around 1-2 % for 1 TeV muons and depends on their η and ϕ .
- Muon p_T resolution : the systematic uncertainty arising due to muon p_T resolution is treated as a shape uncertainty. The muon momentum resolution is smeared by a Gaussian with a width of 1 % in the barrel and 2 % in the endcaps, as recommended by the Muon POG.
- Muon reconstruction efficiency : a momentum and η -dependent uncertainty is considered. The size of the uncertainty is ~ 0.5 % for 1 TeV muons. This is treated as a shape uncertainty.
- Muon trigger: a p_T and η -dependent shape uncertainty is applied. For 1 TeV muons, the magnitude of the uncertainty varies between 3 and 10 % depending on the year and the muon η .
- Electron trigger: like for the muon, a p_T and η -dependent shape uncertainty is applied. For 1 TeV electrons, the trigger efficiency is close to 100 % and thus the uncertainty is low, in the 0.1-0.2 % range.

1. The minimum bias cross section corresponds to the total inelastic proton-proton cross section, cf. chapter 6.

2. For the 2016 $e\mu$ analysis the NNPDF3.0 weights are used since the 2016 $Z' \rightarrow e\mu$ samples were produced with the 3.0 weights.

- Trigger prefire: applied for the data taken in the years 2016-17, when prefire was an issue. The timing of L1 trigger primitives coming out of high η ECAL crystals got progressively misaligned with the rest of the detector, causing a significant fraction of them to be assigned to the previous bunch crossing. Events with high energy deposits in the ECAL $2 < |\eta| < 3$ region would fire the L1 for the wrong bunch crossing, and veto data acquisition for the correct bunch crossing due to trigger rules preventing two consecutive bunch crossings to fire. Prefiring is only a minor issue in this LFV analysis, and the recommended uncertainties are used.
- Electron energy scale: the systematic shape uncertainty on the electron E_T scale is taken into account. The electron energy scale is varied by $\pm 2\%$ [192].
- Electron scale factors: scale factors are applied for electrons that pass the HEEP ID. In 2016, for the barrel, a systematic uncertainty of 1 % is used for E_T below 90 GeV, 3 % for energies higher than 1 TeV, and a linear interpolation from 1 % to 3 % is used for the intermediate range 90-1000 GeV. For the endcap the values are 1 % below 90 GeV, 4 % for energies higher than 300 GeV, and again a linear interpolation from 1 % to 4 % is used for the 90-300 GeV range [192]. In 2017-18, the endcap uncertainties are slightly higher: 2 % at low E_T and 5 % at high E_T , with a linear interpolation between the two. These are shape uncertainties.
- Tau energy scale: the tau p_T and energy are varied by a factor depending on the reconstructed decay mode, the tau p_T , and whether the reconstructed tau is genuine or a muon or electron fake. For taus with $p_T > 100$ GeV, the uncertainty varies between 1.5 and 4 %, depending on the decay mode. This is a shape uncertainty.
- Shape uncertainty on the scale factor for tau reconstruction and identification: a p_T -dependent uncertainty (5 % for 1 TeV taus) is applied.
- Electron-to-tau and muon-to-tau fake rates: η -dependent shape uncertainties are considered. The size of these uncertainties varies from 8-40 % depending on the reconstructed tau η and on whether it is an electron or muon fake. This uncertainty is only applied to simulated events in which a reconstructed tau matches a generated electron or muon, a small proportion of the background.
- MET uncertainties: the MET is the negative \vec{p}_T sum of all objects in a given event. If the energy scale of any object changes, then the MET changes too. The uncertainties on the energy scales of the muons, electrons and taus mentioned above are propagated to the MET. On top of that, uncertainties on the jet energy scale and resolution and their effects on the MET also need to be considered. Finally MET uncertainties arising from other objects (“unclustered energy”) need to be included too.
- Data-driven background: in the $e\mu$ final state a flat 50 % systematic uncertainty is applied, and the same is done in the tau channels. See chapter 8 for a justification of these uncertainties.
- The normalization uncertainties on background cross sections are taken from the scale uncertainties on the cross sections recommended by the CMS Standard Model Physics and Top Physics groups:
 - $t\bar{t}$: 5 %
 - Single top : 5 %
 - Drell Yan : 2 %
 - WW : 3 %
 - WZ : 4 %
 - ZZ : 4 %

- WW shape uncertainty: uncertainties due to missing differential higher order corrections for the WW background are taken into account as the WW background is one the most important analysis backgrounds, especially at high invariant masses. The uncertainty is derived by the differential NLO electroweak corrections to the LO cross section as calculated in [177]. They are computed as a function of the invariant mass of the dilepton system M_{ll} . The correction is parametrized by a second order polynomial with the form:

$$\Delta = 1 - (0.993 - 2.001 \cdot 10^{-4} \times M_{ll} + 2.838 \cdot 10^{-8} \times M_{ll}^2)$$

The nominal WW distributions has to be multiplied by $1 \pm \Delta$ to obtain the up/down shape uncertainty templates. Since this is a symmetric uncertainty, it also covers possible contributions from $\gamma\gamma$ induced processes which are not simulated in the WW simulation samples. This uncertainty corresponds to a 30-50 % variation in the number of WW events for a dilepton mass scale of 2 TeV.

- Top shape uncertainty: during LHC Run 1 and Run 2 it was found that the p_T spectra of top quarks in data were significantly softer than those predicted by the various NLO simulations we use for this analysis. In other words, there are more events than expected for low p_T tops but fewer events than expected for high p_T tops. The effect is more enhanced in the tail of the p_T spectrum. We have come up with a function for reweighting our POWHEG NLO $t\bar{t}$ sample to the most precise available theory prediction (NNLO QCD + NLO EW) [178] (see chapter 7). The uncertainty on this reweighting is the main uncertainty for $t\bar{t}$ at high mass. As shown in figure 9.1, two different systematics are extracted, one related to the choice of PDFs, and the other to the factorization scale used by the event generator (Qscale). This uncertainty has a similar size to the WW shape uncertainty, corresponding to a 30-50 % variation in the number of $t\bar{t}$ events for a dilepton mass scale of 2 TeV.

All uncertainties are considered to be correlated across the different data-taking years, except when the relevant CMS study group has determined that an uncertainty is uncorrelated. No correlation between the different final states is considered, the analysis results are presented independently for each of the final states, with no combination.

Table 9.1 summarizes the basic parameters of all systematic uncertainties. Figure 9.2 shows the relative importance of the various systematic uncertainties in the $\mu\tau$ final state, for background processes and one signal sample. At high masses, the dominating background uncertainties are the muon systematics, the statistical uncertainty and the uncertainty on the WW PDF. For the Z' signal, statistical and muon-related uncertainties also dominate, and the signal PDF uncertainty is the third most important contribution.

9.2 Lepton kinematics and mass plots

In this section, data events passing signal region selections are compared to background estimations coming from simulations and data-driven methods. Data from the full 2016-18 data-taking period are shown, and systematic uncertainties on the backgrounds are displayed to aid the comparison.

Figure 9.3 shows lepton kinematical variables in the $e\mu$ final state. There is good agreement between observed and expected events. The HEEP electron η distributions shows two holes at values corresponding to the crack regions between the ECAL barrel and endcaps. These regions are vetoed due to the poor electron reconstruction there.




Figure 9.1 – PDF and Qscale systematic uncertainties on the simulated $t\bar{t}$ background, as a function of the dilepton mass $M_{\ell\ell}$. Since there are two tops in each event, a geometric average of the uncertainties calculated for each top p_T is used.

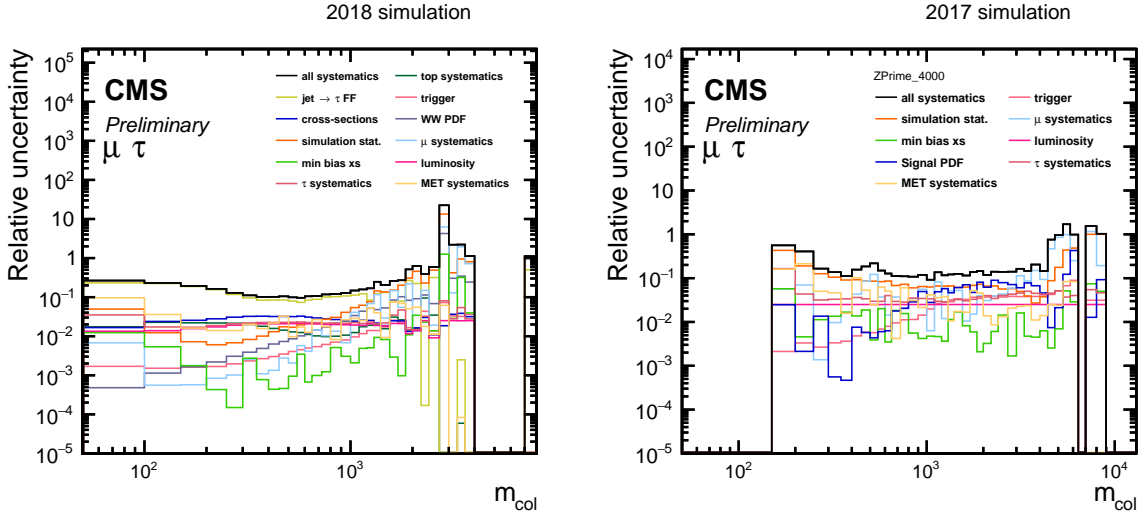


Figure 9.2 – Relative uncertainties in the $\mu\tau$ channel. (Left) Relative uncertainties on the background, (right) relative uncertainties on the 4 TeV Z' signal. The left plot was constructed by dividing the size of each uncertainty by the total expected background, in the right plot the uncertainties were divided by the expected yield of the signal sample. Uncertainties are grouped together to reduce clutter, e.g. all uncertainties applying to muon scale factors, energy scale corrections or p_T resolution are grouped under “ μ systematics”.

Systematic	Type	Channels	Correlated across years?
Trigger	Shape	all	no
PU reweighting	Shape	all	yes
Luminosity	Normalization	all	no
Prefiring	Shape	all	yes
$t\bar{t}$ cross section	Normalization	all	yes
WW cross section	Normalization	all	yes
ST cross section	Normalization	all	yes
DY cross section	Normalization	all	yes
WZ, ZZ cross section	Normalization	all	yes
jet \rightarrow e, jet \rightarrow μ fakes	Normalization	$e\mu$	yes
jet \rightarrow τ fakes	Normalization	$e\tau, \mu\tau$	yes
Muon p_T scale	Shape	$e\mu, \mu\tau$	yes
Muon resolution	Shape	$e\mu, \mu\tau$	no
Muon reconstruction	Shape	$e\mu, \mu\tau$	no
Muon isolation	Shape	$e\mu, \mu\tau$	no
Ele E_T scale	Shape	$e\mu, e\tau$	no
Ele HEEP ID	Shape	$e\mu, e\tau$	yes
Tau p_T scale	Shape	$e\tau, \mu\tau$	no
$e \rightarrow \tau$ fake rate	Shape	$e\tau, \mu\tau$	no
$\mu \rightarrow \tau$ fake rate	Shape	$e\tau, \mu\tau$	no
DeepTau ID scale factor	Shape	$e\tau, \mu\tau$	no
MET	Shape	$e\tau, \mu\tau$	no
WW shape	Shape	all	yes
$t\bar{t}$ shape	Shape	all	yes
Signal PDF	Shape	all	yes

Table 9.1 – Systematic uncertainties considered in this LFV data analysis. Normalization uncertainties are modeled with log-normal distributions, while shape uncertainties are extrapolated from two histograms representing $\pm 1 \sigma$ variations. Some uncertainties are considered to correlated across the different years.

Some control variables in the $e\tau$ channel are shown in figure 9.4. Electrons and taus in the barrel/endcap ECAL transition are vetoed, hence the fall in the number of events for $|\eta| \sim 1.5$ for both leptons (the histogram binning is not fine enough to show zero event). The MET plot shows excellent agreement. When all systematic uncertainties are taken into account, there is overall good agreement between data and background expectations.

Figure 9.5 shows some control variables for the $\mu\tau$ channel. The m_T variable has been replaced by the relative sign between the two leptons, to show that same sign events are almost exclusively jet $\rightarrow \tau$ fakes. Just as in the $e\tau$ channel, there is an overall good agreement between observations and background expectations. The expected lack of events for tau $|\eta| \sim 1.5$ can be observed, and the MET variable is under control with background expectations describing the data with high accuracy.

With no disagreement between data and background estimations appearing in the plots of the control variables, it is no surprise that the mass plots in figure 9.6 show no excess in data events with respect to Standard Model expectations. The number of high mass jet $\rightarrow \tau$ fakes is substantially higher in the $e\tau$ channel compared with the $\mu\tau$ channel, this is due to the higher likelihood of a jet faking an electron compared to that of a jet faking a muon. Details are provided in appendix C.

Since no new physics has been found, a statistical analysis is performed to determine which versions of the signal models considered in this analysis (LFV Z' , QBH, RPV SUSY) are ruled out.

9.3 Statistical analysis and limits on new physics models

No significant excess of data events over the expected background is observed. An upper limit at the 95 % confidence level (CL) on the cross section \times branching ratio (σB) is determined using a Bayesian binned-likelihood approach, assuming a uniform prior for the signal cross section. The advantage of using a uniform prior is that, in the absence of background events, the upper limit on a Poisson process obtained by Bayesian methods is identical to the upper limit obtained by frequentist methods [193]. In the presence of background, that coincidence no longer holds, but the Bayesian limit is more conservative [194]. The upper limit on the new physics cross section is directly related to the number of potential signal events observed, which follows a Poisson distribution.

The nuisance parameters associated with the systematic uncertainties are modeled through log-normal distributions for uncertainties in the normalization. Uncertainties in the shape of the distributions are modeled through template morphing techniques, i.e. shapes representing one standard deviation fluctuations in opposite directions are provided, and the fitting tool extrapolates an uncertainty distribution from them. A Markov Chain Monte Carlo method is used for integration. After integrating over all nuisance parameters for each mass hypothesis, the posterior probability density is calculated as a function of σB for yields at the 95 % CL upper limit, for all three signal models: LFV Z' , QBH and RPV SUSY. The fitting tool used for this statistical analysis is the Higgs Combine tool [195].

The analysis focusing on specific signals, as described in the previous paragraph, assumes a certain signal shape in the discriminating variables, i.e. the collinear mass in the tau channels and the invariant mass in the $e\mu$ channel. However, alternative and unforeseen new physics processes yielding lepton flavor violating final states could cause a data excess with a different

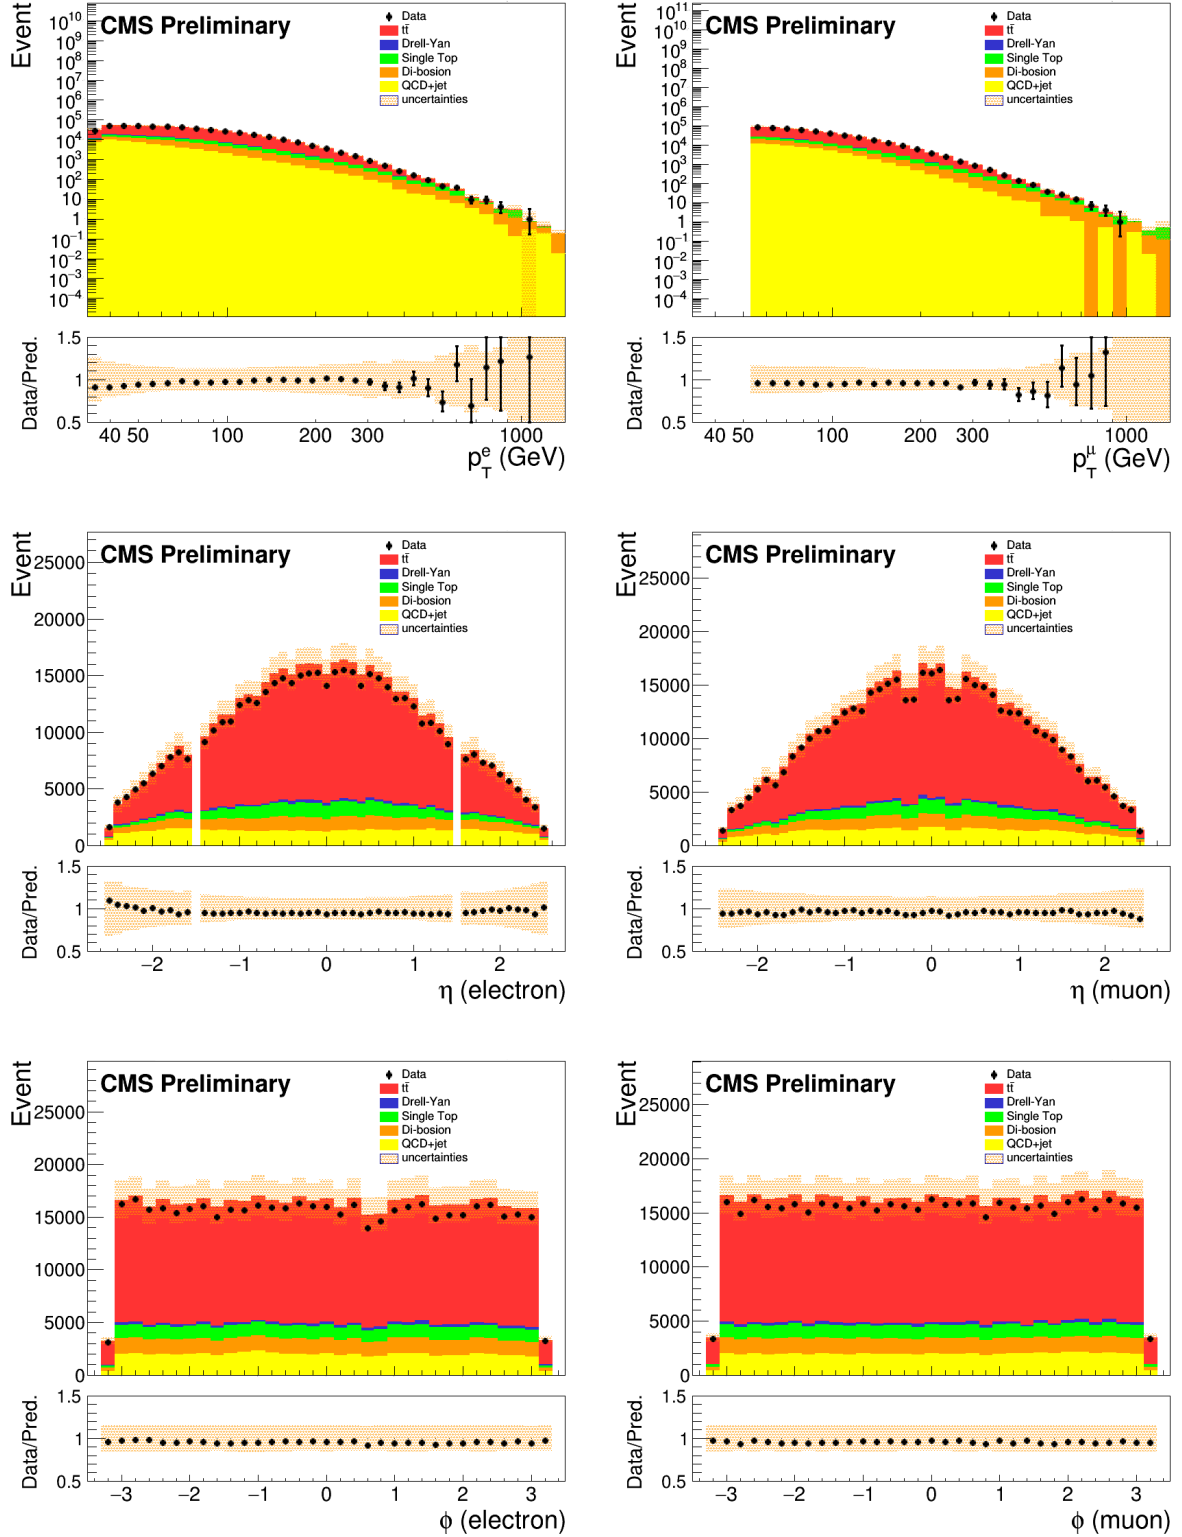


Figure 9.3 – Control variables in the $e\mu$ channel, with full Run 2 data. From left to right, top to bottom: (1) $p_T(e)$, (2) $p_T(\mu)$, (3) $\eta(e)$, (4) $\eta(\mu)$, (5) $\phi(e)$, (6) $\phi(\mu)$. The shaded uncertainty band represents both statistical and systematic uncertainties.




Figure 9.4 – Control variables in the $e\tau$ channel, with full Run 2 data. From left to right, top to bottom: (1) $p_T(e)$, (2) $p_T(\tau)$, (3) $\eta(e)$, (4) $\eta(\tau)$, (5) $m_T(e, \cancel{p}_T)$, (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

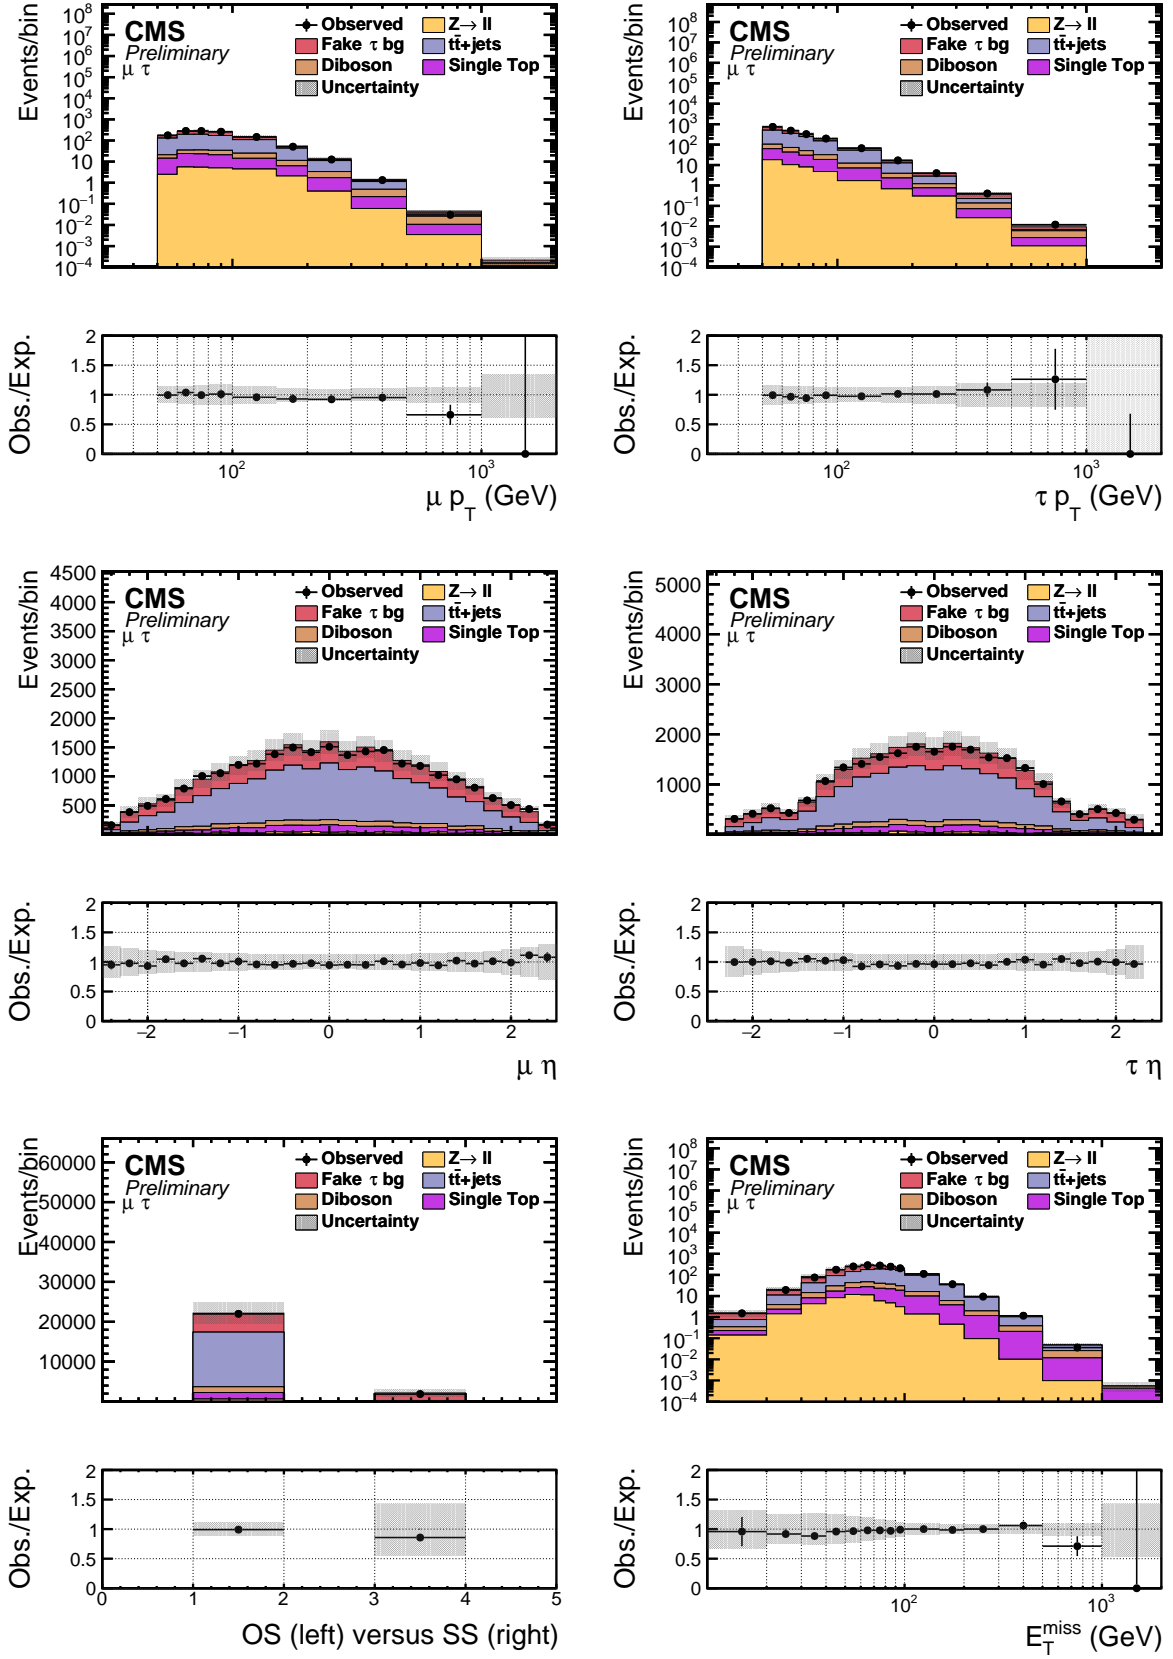


Figure 9.5 – Control variables in the $\mu\tau$ channel, with full Run 2 data. From left to right, top to bottom: (1) $p_T(\mu)$, (2) $p_T(\tau)$, (3) $\eta(\mu)$, (4) $\eta(\tau)$, (5) relative sign of $\mu\tau$, (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

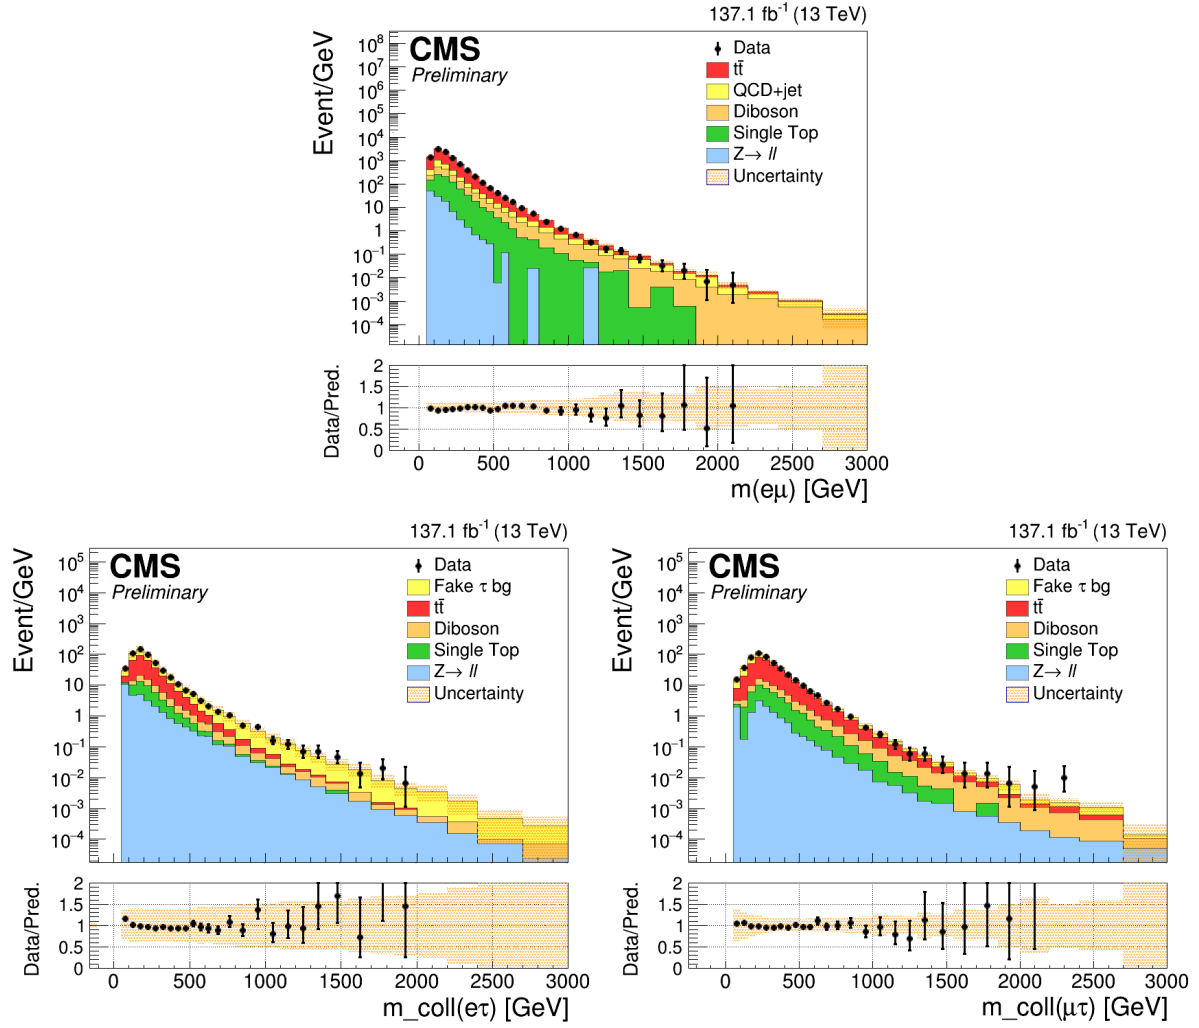


Figure 9.6 – Comparison of data and expected background mass distributions. The $e\mu$ invariant mass is shown on top, the $e\tau$ collinear mass is shown on the bottom left and the $\mu\tau$ collinear mass on the bottom right. The shaded uncertainty band represents both statistical and systematic uncertainties. No significant data excess above background expectations is observed.

shape. It is thus interesting to determine a model-independent cross section limit, by using a single bin ranging from a minimum threshold on collinear mass or invariant mass up to infinity. No assumption on the shape of the signal distribution has to be made other than that of a flat product of acceptance times efficiency, $A\epsilon$, as a function of the mass variable. In order to determine the limit for a specific model from the model-independent limit shown here, only the model-dependent part of the efficiency needs to be applied. The experimental efficiencies for signal events are already taken into account.

A factor f_m that reflects the effect of the threshold mass m^{min} on the signal efficiency is determined by counting the events with mass $m > m^{min}$ and dividing the result by the total number of generated events. Detector effects on the event reconstruction efficiency are nearly constant over the entire mass range probed here, therefore f_m can be evaluated at generator level. A limit on the product of the cross section, branching fraction, acceptance and efficiency $(\sigma BA\epsilon)_{excl}$ can be obtained by dividing the excluded cross section of the model-independent limit $(\sigma BA\epsilon)_{MI}$ (which is a function of m^{min}) by the calculated fraction $f_m(m^{min})$:

$$(\sigma BA\epsilon)_{excl} = \frac{(\sigma BA\epsilon)_{MI}(m^{min})}{f_m(m^{min})} \quad (9.1)$$

Here, B is the branching fraction of the new particle decaying to an LFV final state. Models with a theoretical cross section $(\sigma B)_{theo}$ larger than $(\sigma B)_{excl}$ can be excluded. The procedure described here can be applied to all models involving the two-body decay of a massive state, which exhibit back-to-back kinematics similar to those of a generic Z' . But it is also possible to obtain excluded cross sections for models with different kinematic properties, in that case the fraction of events $f_m(m^{min})$ must be determined for the particular model considered.

Figure 9.7 shows the exclusion plots at the 95 % confidence level of the LFV Z' and RPV SUSY models for all three final states. Figure 9.8 shows the exclusion plots for the QBH model and the model-independent limits. 95 % CL limits on the resonance mass or threshold mass of the three specific models considered can be extracted from the plots, and are summarized in table 9.2. Details on the signal models are given in chapter 6. Both observed and expected limits are shown: observed limits take into account data events recorded in the signal region, while expected limits show what the result would be if the recorded data had corresponded exactly to background expectations. Expected limits thus remove effects from possibly spurious upwards or downwards fluctuations of the data, and represent instead only the experimental sensitivity to new physics of the analysis. Whenever possible, the results from this analysis are compared to other CMS and ATLAS results, showing substantial improvement. It should be stressed in particular that the searches in the tau channels had never been done before in the CMS collaboration.




Figure 9.7 – Expected (black dashed line) and observed (black solid line) 95% CL upper limits on the product of cross section and branching fraction, using full LHC Run 2 data, for a Z' boson with LFV decays (left column) and a $\tilde{\nu}_\tau$ derived from RPV SUSY (right column). Two theoretical curves are shown for the RPV model, one with $\lambda = \lambda' = 0.1$ and another with $\lambda = \lambda' = 0.01$. The top row shows results for the $e\mu$ channel, the middle row for the $e\tau$ channel and the bottom row for the $\mu\tau$ channel. The shaded bands represent the one and two standard deviation (s.d.) uncertainty bands.

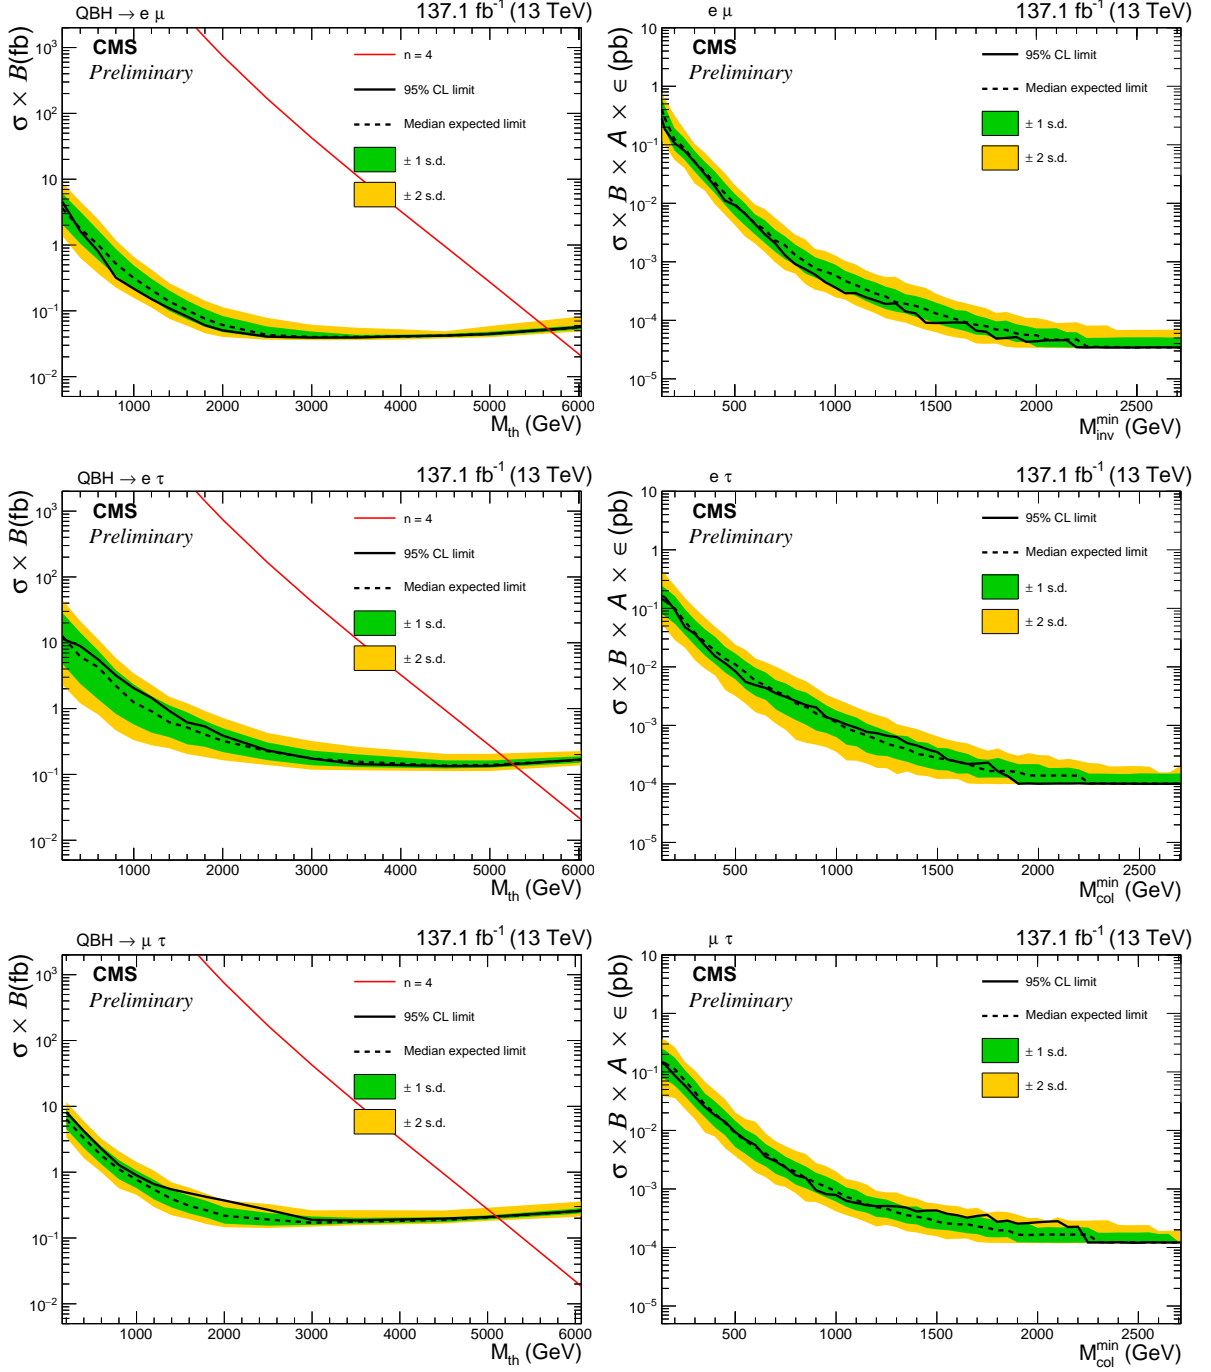


Figure 9.8 – Expected (black dashed line) and observed (black solid line) 95% CL upper limits on the product of cross section and branching fraction for an ADD $n=4$ quantum black hole (left column). Model independent upper limits on the product of cross section, branching fraction, acceptance and efficiency are shown in the right column. The top row shows results for the $e\mu$ channel, the middle row for the $e\tau$ channel and the bottom row for the $\mu\tau$ channel. The shaded bands represent the one and two standard deviation (s.d.) uncertainty bands.

		Observed (expected) limits in TeV		
		$e\mu$	$e\tau$	$\mu\tau$
LFV Z'	This analysis	5.0 (4.9)	4.2 (4.2)	4.1 (4.3)
	2016 CMS	4.4 (4.4)	—	—
	2016 ATLAS	4.5 (4.3)	3.7 (3.7)	3.5 (3.5)
RPV SUSY $\lambda = \lambda' = 0.1$	This analysis	4.2 (4.2)	3.7 (3.7)	3.7 (3.7)
	2016 CMS	3.8 (3.8)	—	—
RPV SUSY $\lambda = \lambda' = 0.01$	This analysis	2.3 (2.3)	1.0 (1.2)	1.7 (1.7)
	2016 CMS	1.7 (1.9)	—	—
QBH, ADD n=4	This analysis	5.7 (5.7)	5.3 (5.3)	5.1 (5.1)
	2016 CMS	5.3 (5.3)	—	—

Table 9.2 – Limits on new physics at 95 % CL obtained in the LFV analysis described in this thesis are presented and compared to results from the previous CMS analysis. For the benchmark Z' model, results can be directly compared to the latest ATLAS analysis too, but unfortunately no such comparison can be made for other signals as different versions of those models were considered in the ATLAS search. The results from the ATLAS analysis [81] and the previous CMS analysis [82] were obtained with 2016 LHC data. Since this is the first time such an analysis has been performed with the CMS detector in the tau channels, tau channel results are not available for the 2016 CMS analysis. For the LFV Z' and the RPV SUSY models, the limits correspond to the minimum resonance masses (of the Z' boson or the tau sneutrino $\tilde{\nu}_\tau$) that cannot be excluded by the data analysis. In the RPV model, the only non-zero RPV couplings are those coupling the $\tilde{\nu}_\tau$ to down-type quarks (λ') and allowing its decay to a specific LFV final state (λ). For the QBH model (ADD model with n=4 extra spatial dimensions), the limits correspond to the minimum unexcluded threshold masses.

9.4 Discussion of the results and prospects

We can observe that in all cases the $e\mu$ limits are stronger than those from the tau channels. This is due to the fact that the tau channels have worse signal acceptance than the $e\mu$ channel since only hadronically decaying taus are identified, while taus with fully leptonic decays are ignored. The hadronic tau identification efficiency is also lower than that of high energy electrons or muons. Limits in both tau channels are comparable to each other for the Z' model and the high coupling ($\lambda = \lambda' = 0.1$) RPV SUSY model. However, the $\mu\tau$ limit is stronger than the $e\tau$ limit for the low coupling RPV model ($\lambda = \lambda' = 0.01$), and the opposite is true for the QBH limits. Background uncertainties play a more important role at low masses, where the $\text{jet} \rightarrow \tau$ fakes contribution is larger in the $e\tau$ channel than in the $\mu\tau$ channel. A large uncertainty is assigned to this background, making the $e\tau$ limit considerably less strong than the $\mu\tau$ one. However, at higher masses, signal acceptance and efficiency become more important than background uncertainties. Muon efficiency decreases for increasing p_T and is lower than that of high energy electrons, making the $e\tau$ limit stronger than the $\mu\tau$ one for the QBH model, where the limits are above 5 TeV for all channels.

As shown in figure 9.9 the acceptance \times efficiency values for the QBH and RPV signals are comparable over the whole mass range, and significantly higher than for the Z' signal. The Z' resonance is spin-1, while the QBH and the RPV $\tilde{\nu}_\tau$ are spin-0. This causes the η distributions of the final state leptons to be wider for the Z' signal: there are relatively more events with high muon or tau eta, falling outside of the detector acceptance (cf. figure 9.10). This difference in acceptance is part of the reason why the QBH limits are stronger than the Z' limits. The RPV limits are weaker because the Z' couplings are larger (they are assumed to be unity in the

benchmark model).

The RPV $\tilde{\nu}_\tau$ and Z' resonances can be seen as two extreme cases between which resonances from other models lie. The $\tilde{\nu}_\tau$ has a very small width, with the width of the reconstructed resonance determined by experimental effects, while the 3 % theoretical width of the Z' is large and visible in the reconstructed mass distribution. The expected and observed limits on the product of cross section and branching fraction given in figures 9.7- 9.8 are still valid for lower values of the coupling constants, as long as the signal acceptance and efficiency remain the same. This assumption is valid for the SUSY model where changes in the coupling constants are expected to change the signal yield but not its shape (as long as the coupling is not too large), and this has indeed been verified for the two cases presented here: $\lambda = \lambda' = 0.1$ and $\lambda = \lambda' = 0.01$. For the Z' signal, however, the resonance width is expected to change with the coupling, with non-trivial effects on the signal acceptance and efficiency. The generation of different signal samples would be necessary to check the magnitude of these effects.

These results can also be compared to other analyses:

- **$Z' \rightarrow \tau\tau$ searches** consider similar final states, since tau leptons can decay to electrons or muons in addition to hadrons. However, there is an important difference: since one or two neutrinos are produced in each tau decay, there is always MET in a $\tau\tau$ event, and this MET is not expected to be aligned with the hadronic tau in a $\mu\tau_h$ or $e\tau_h$ event. The mass variables considered here, the visible $e\mu$ invariant mass and the collinear mass in the tau channels, are not appropriate for such a search. However, the ditau searches are still a useful comparison since new resonances are expected to have flavor-diagonal as well as LFV couplings. The latest lower limits on the mass of a Sequential Standard Model (SSM) Z' are 2.1 TeV for CMS [196] (with 2015 data) and 2.4 TeV for ATLAS [197] (with 2015-16 data). The LFV limits are stronger, the ditau limits are only competitive if LFV couplings are heavily suppressed compared to flavor-diagonal couplings.
- **$Z' \rightarrow ll$ ($l=e, \mu$) searches** are also an important consideration, the flavor-diagonal couplings of new resonances can be ee and $\mu\mu$, and not just $\tau\tau$. The $Z' \rightarrow ee$ and $Z' \rightarrow \mu\mu$ searches have a comparable mass reach to the $e\mu$ channel in this analysis, the CMS observed limits for the SSM Z' with full Run 2 data [186] are respectively 4.7 and 4.9 TeV for the dielectron and dimuon channels, to be compared with the 5.0 TeV of the $e\mu$ channel. The ATLAS limits [198] are similar, stronger in the dielectron channel (4.9 TeV) but weaker in the dimuon channel (4.5 TeV). However, for the tau LFV channels to be competitive, their branching ratios would need to be about 3-4 times higher than those of the non-tau channels.
- **LFV decays of muons or taus.** Let us consider $\mu \rightarrow e\gamma$ and $\tau \rightarrow l\gamma$. This paper [74] argues that an upper limit of order 10^{-13} (such as the one obtained by the MEG experiment [63]) on the $\mu \rightarrow e\gamma$ branching ratio would translate to a lower limit of order 1000 TeV on the mass scale of LFV new physics (see also figure 9.11), by assuming as a benchmark that Z' LFV couplings are unity. The $l \rightarrow l'\gamma$ branching ratios are proportional to $\frac{1}{M_{Z'}^4}$ where $M_{Z'}$ is the mass scale of the LFV new physics [199, 200]. So for $\tau \rightarrow l\gamma$ decays, where the branching ratio limits are of order 10^{-8} [65, 66], this would translate to a mass scale of ~ 10 -100 TeV. For LFV new physics involving $e\mu$ mixing, the low energy constraints are stronger, while for LFV new physics involving taus, the low and high energy constraints are comparable, though the low energy limits look slightly stronger. Nevertheless, the high energy results are still an interesting independent result, they are competitive for lower coupling values and there are models for e.g. quantum black holes in which low energy phenomena are heavily suppressed (cf. chapter 2).




Figure 9.9 – Acceptance \times efficiency in the $\mu\tau$ channel for different mass points (resonance mass or threshold mass) of each of the three signals.

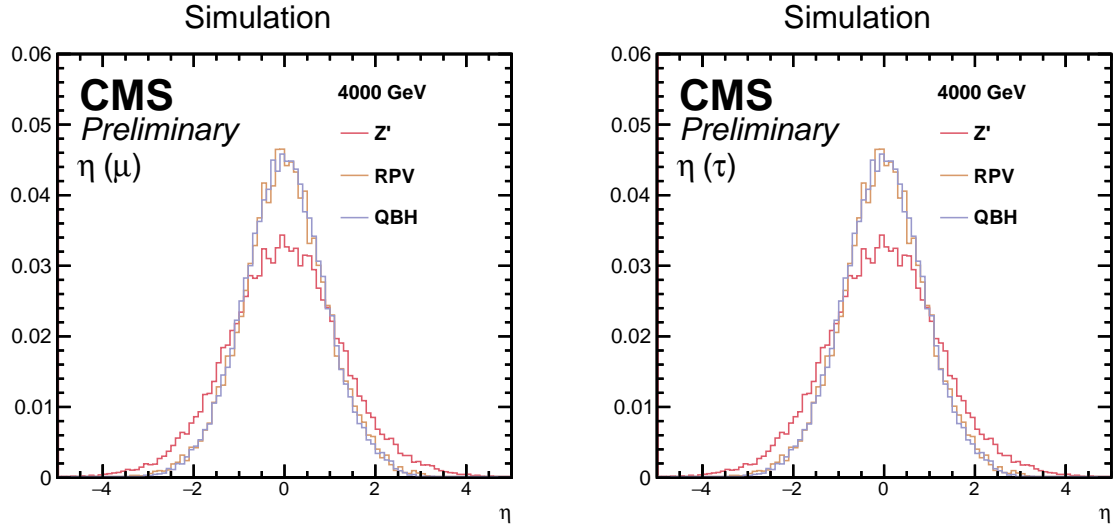


Figure 9.10 – Normalized pseudo-rapidity distributions at the generated level for the three types of signal, at the 4 TeV mass point, in the $\mu\tau$ channel. (Left) Muon η distributions, (right) visible hadronic tau η distributions. Pseudo-rapidity acceptance is $|\eta| < 2.4$ for muons and $|\eta| < 2.3$ for hadronic taus.

Prospects for improvement

This is the first CMS analysis looking at lepton flavor violation with tau leptons at the high mass frontier. The results are novel and state of the art, but there is room for improvement in future analyses:

- Trigger efficiency can be improved in the $\mu\tau$ channel, it is currently only $\sim 90\%$ for high p_T muons (see figure 7.2). A more efficient trigger combination could be used, or better performing single muon triggers developed.
- A better estimation of the signal PDF uncertainties would be possible with new signal samples generated with all the necessary PDF weights. To estimate the Z' PDF uncertainties, it was necessary to reweight the samples, which introduces some biases in the signal yields. For the other signals, we relied on work done for the 2016 $e\mu$ analysis [82].
- The WW PDF uncertainty could also be more precisely estimated, by doing a procedure similar to the one used for the $t\bar{t}$ background and reweighting the samples to NNLO QCD + NLO EW precision, and then using the differences between NNLO QCD + NLO EW and NLO predictions to determine uncertainties.
- Signal acceptance in the tau channels can be improved by including the fully leptonic decays of the tau and not just the hadronic decays. Like in the LFV H boson searches [75, 76], the addition of the $\mu\tau_e$ and $e\tau_\mu$ channels would be possible.
- Another possibility for improving the signal acceptance in the tau channels would be to use a looser working point of the DeepTau anti-jet discriminator. The uncertainty on the jet $\rightarrow \tau$ fakes would need to be better understood. The dependence of the fake rates on jet flavor is a promising start, but more work would be needed to better understand the flavor composition of this background.
- Statistical uncertainties are still high, so more data, such as that collected during future HL-LHC runs, would help.
- A combined fit of the three channels could be performed to provide stronger limits for specific models predicting the relative yields of the different LFV final states³.

The constraints on lower mass LFV resonances could be made stronger, but the analysis selections would need a higher emphasis on background rejection. In the $e\mu$ channel, events with high MET could be rejected. In the tau channels, the MET and tau p_T could be required to be aligned, and the m_T selection might need to be dropped in order not to reject low mass signal events. A better understanding of the uncertainties on the fake jet backgrounds would have a high impact on low mass limits, especially in the $e\tau$ channel.

3. e.g. flavor-democratic models, though in this case the contribution of the tau channels would be small compared to that of the $e\mu$ channel.




Figure 9.11 – Sensitivity of $\mu \rightarrow e$ and $\mu \rightarrow e\gamma$ experiments to the LFV new physics mass scale Λ , as a function of the κ parameter in the effective LFV Lagrangian [74] (see also chapter 1).

Chapter summary

In order to properly interpret the comparison between data and Standard Model background expectations in the signal region, all sources of uncertainty, both statistical and systematic, need to be taken into account. Some systematic uncertainties relate to various biases that might affect detector response, for example the identification of particles and the calibration of their energy. Most of these uncertainties are standard recommendations given by particle object groups of CMS after a thorough study of the relevant issues. Other uncertainties are theoretical, such as that on the $t\bar{t}$ cross section, and derived from published papers. Finally, some of these uncertainties, such as those on the fake lepton backgrounds, were determined by original work done specifically for the LFV analysis. Comparisons between data and SM expectations show agreement in the signal region, after all systematic effects are taken into account. No localized excess of data is observed in the plots of the mass variables of each channel, used to discriminate signal from background. A statistical analysis is then performed in order to obtain 95 % confidence level exclusion plots for the three specific models considered: RPV SUSY with two non-vanishing RPV couplings relating to tau sneutrino production and decay, an ADD model with four extra spatial dimensions, and a Z' model where the Z' has exclusively LFV decay modes. The observed lower limit on the LFV Z' mass is respectively 5.0, 4.2 and 4.1 TeV for the $e\mu$, $e\tau$ and $\mu\tau$ channels. Model-independent limits were also obtained. The results are a substantial improvement over previous CMS and ATLAS analysis, and the tau channel searches in particular are entirely novel in the CMS collaboration. The chapter concluded with a discussion of how these results compare with other high-energy and low-energy searches, and a review of possible improvements in future versions of this analysis.

My personal contributions

My contributions to the production of the signal region distributions and the statistical analysis are the following:

- I designed and ran the $\mu\tau$ analysis code and produced the signal region plots in this channel.
- Occasionally I ran my own version of the $e\tau$ analysis code to cross-check results if any discrepancies were observed.
- I studied different possible ways of estimating the signal PDF uncertainties, since unfortunately the simulation samples we had did not contain the MC replicas which are used in the standard method for estimating these uncertainties.
- In general, I participated in our regular meetings about the analysis strategy, where we decided which selections to use and how to estimate systematic uncertainties. As the contact person for the internal CMS review of the analysis, I had to be up-to-date with all the aspects of the analysis, in all three channels.

Conclusion

This thesis presents a search for new physics phenomena in lepton flavor violating final states. The Standard Model of particle physics implies that, when the neutrino masses are set to zero, the electron, muon and tau numbers are all separately conserved. However, small changes to the SM generically lead to the non-conservation of lepton flavor, which means that the search for such phenomena is sensitive to a large range of new physics models. The SM, while very successful in explaining particle physics phenomena, still has a number of outstanding issues that are addressed by new physics models. Lepton flavor violation has already been observed in the neutrino sector (implying that neutrinos have non-zero masses) but never in charged leptons. Searches for LFV phenomena have already been performed in lower energy experiments and in the LHC itself. The work presented in this thesis is novel for the tau channels using CMS data, and for all three channels the results improve on previous results from ATLAS and CMS.

The data used in the analysis was collected by CMS during LHC Run 2, in the years 2016-18. It corresponds to an integrated luminosity of 137.1 fb^{-1} of proton-proton collisions, at a center-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. Datasets constructed with single muon, single electron and single photon triggers are used. Collision events are reconstructed with CMS software, with taus only reconstructed when their decays are hadronic, and electrons and muons using identification algorithms dedicated to the high energies and momenta expected in signal events. Event selection is kept simple, in order to make the analysis as model-independent as possible. Events selected need to have two different-flavor, well-identified and well separated high momentum leptons in the detector acceptance. In the tau channels, a few extra selections are added: a transverse mass requirement to reduce backgrounds with wrongly-identified taus, and an extra lepton veto to remove overlap with other final states in the analysis and also to reduce backgrounds with same-flavor leptons. Background estimation relies on simulations when both leptons are real, but data-driven methods were used to estimate events with $\text{jet} \rightarrow e$ and $\text{jet} \rightarrow \mu$ fakes in the $e\mu$ channel, and events with $\text{jet} \rightarrow \tau$ fakes in the tau channels. The simulations are reweighted by scale factors determined by experimental calibration or theoretical corrections.

The discriminating variables in the analysis are the $e\mu$ invariant mass, and the collinear mass m_{col} in the tau channels, incorporating information about the MET. A localized excess of data events in these distributions is the expected signal signature. However, in all three final states, no excess in data events relative to expectations is observed. As a consequence, a statistical analysis is performed in order to find upper limits at the 95 % confidence level on the cross sections of new physics processes. A comprehensive list of systematic uncertainties is taken into account, including theoretical uncertainties and detector calibration and other

experimental effects. The lower limits on the LFV Z' resonance mass are respectively 5.0, 4.2 and 4.1 TeV for the $e\mu$, $e\tau$ and $\mu\tau$ channels. For the R-parity violating SUSY-inspired tau sneutrino with couplings $\lambda = \lambda' = 0.1$, the limits are respectively 4.2, 3.7 and 3.7 TeV. For quantum black holes derived from an ADD model with four large extra dimensions, the lower limits on the threshold mass are respectively 5.7, 5.3 and 5.1 TeV.

My personal contributions to the work presented in this thesis can be classified into three categories: work on the HLT, work in tau studies, and work in the $\mu\tau$ channel of the analysis. My HLT work was to develop a new tool for the estimation of trigger rates, and using that tool to provide feedback on newly-deployed or soon-to-be-deployed HLT menus. I also studied the overlap in the rates of different CMS datasets, providing recommendations which improved the way CMS data was collected in 2018. For the TAU particle object group, I studied the efficiency of hadronic tau reconstruction and identification in 2017 data, though I did not provide the final scale factor recommendations. My work in the analysis proper was mostly focused in the $\mu\tau$ channels, making the signal region and control region distributions, and estimating the $\text{jet} \rightarrow \tau$ scale factors, and the uncertainty on them. I did the work showing why the collinear mass is a better discriminating variable for the tau channels than the visible mass or the total mass. I also tuned the definition of the collinear mass so that it could handle events with misaligned tau and MET in a consistent way. I also determined the best way to evaluate signal PDF uncertainties with the samples we had available, and I participated in the discussions about general analysis strategy: what the selections should be, what systematic uncertainties need to be taken into account. And I closely collaborated with my colleague working in the $e\tau$ channel, cross-checking any discrepancies between our results, and comparing numbers such as the $\text{jet} \rightarrow \tau$ fake factors and the signal efficiencies.

The results described in this thesis represent the new experimental state of the art in collider searches at the energy frontier for lepton flavor violating physics. They complement low energy LFV searches as well as high energy same flavor dilepton searches for new resonances. Much of the LFV high-energy parameter space has already been excluded by the present analysis, but improvements can be expected by repeating this analysis with higher signal efficiency and with the larger amounts of data expected during the High Luminosity LHC runs in the coming decades. The increase in LHC center-of-mass energy to its nominal 14 TeV will also somewhat increase the sensitivity of future analyses to new physics.

LHC searches optimized for lower mass LFV resonances would provide interesting complementary results. The fake jet background uncertainties would need to be better understood, perhaps with a deeper exploration of the fake rate dependence on jet flavor. The selections would need to be adapted, notably by placing a higher emphasis on background rejection.

Observation of LFV phenomena may also come from the many lower-energy experiments focusing on tau decays and $\mu \rightarrow e$ processes (either decay or conversion) that will come online in the near future. In many respects the low energy searches are already sensitive to higher mass scales than the searches at the energy frontier, especially for $e\text{-}\mu$ mixing, though the high mass tau channels remain competitive for lower values of the Z' coupling constant. The devil is in the details of the precise hierarchy of the diagonal and off-diagonal leptonic couplings of the new resonances.

Special relativity and particle physics conventions

Lorentz boosts

In special relativity, Lorentz boosts relate two inertial frames of reference which are moving at a constant velocity relative to one another. Assuming the difference in velocity between the two frames is along the z -direction and has a magnitude v , one can write the new coordinates after a Lorentz boost as:

$$\beta = \frac{v}{c}, \quad \gamma = \frac{1}{\sqrt{1 - \beta^2}}$$

$$\begin{cases} z' = \gamma(z - \beta ct) \\ ct' = \gamma(-\beta z + ct) \end{cases}$$

where t is the time coordinate. γ is called the Lorentz factor and diverges as v approaches c , the speed of light in a vacuum. The energy E and momentum p_z transform similarly:

$$\begin{cases} \frac{E'}{c} = \gamma\left(\frac{E}{c} - \beta p_z\right) \\ p'_z = \gamma\left(-\beta \frac{E}{c} + p_z\right) \end{cases}$$

Natural units

Given the ubiquity of the constant c factor in particle physics, it is often set to 1 to simplify equations. With a careful dimensional analysis, it is always possible to add the missing c factors back in to express quantities in SI units. For the same reasons, the Planck constant \hbar is often set to 1 as well. In the so-called natural units of high energy physics, $c = \hbar = 1$, and energies, momenta and masses are expressed in electron-Volts (eV) or multiples thereof. These units are used throughout this thesis.

Sum over repeated indices

Another way of simplifying equations is to use the convention that repeated indices imply a sum over these indices. For example, when contracting two Lorentz 4-vectors a^μ and b^μ , instead of writing

$$\sum_{\mu=0}^3 a_\mu b^\mu$$

one can simply write

$$a_\mu b^\mu$$

with the sum left implicit.

Particles and antiparticles

In the CMS experiment, particles and their respective antiparticles are detected using the same subdetectors, and reconstructed with the same algorithms. Other than the sign of their charge, and as a consequence the direction in which their tracks are bent by the magnetic field, muons and antimuons look the same to CMS. This is why it is common in high energy physics to lump together particles and their antiparticles. In this thesis, electrons and positrons are often commonly referred to as “electrons”, neutrinos and antineutrinos as “neutrinos”, etc.

Rapidity and pseudo-rapidity

Before explaining pseudo-rapidity, a description of the related concept of rapidity is necessary. If a particle has energy E and momentum along the z-axis p_z its rapidity y is given by the following formula:

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right)$$

Lorentz boosts along the z-axis preserve differences in rapidity:

$$\begin{aligned} y' = \frac{1}{2} \ln \left(\frac{E' + p'_z}{E' - p'_z} \right) &= -\frac{1}{2} \ln \left(\frac{\gamma(E - \beta p_z) + \gamma(-\beta E + p_z)}{\gamma(E - \beta p_z) - \gamma(-\beta E + p_z)} \right) \\ &= \frac{1}{2} \ln \left(\frac{(1 - \beta)(E + p_z)}{(1 + \beta)(E - p_z)} \right) \\ &= y + \frac{1}{2} \ln \left(\frac{1 - \beta}{1 + \beta} \right) \end{aligned}$$

The rapidity is shifted by a constant factor no matter the energy or momentum of the particle, which means that Lorentz boosts do conserve differences in rapidity. If θ is the polar angle of a particle’s momentum, its pseudo-rapidity is defined as:

$$\eta = -\ln \tan \left(\frac{\theta}{2} \right) \tag{A.1}$$

For massless particles, rapidity y and pseudo-rapidity η coincide:

$$\begin{aligned}
y &= \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) = -\frac{1}{2} \ln \left(\frac{|\vec{p}| + p_z}{|\vec{p}| - p_z} \right) \\
&= \frac{1}{2} \ln \left(\frac{1 + \frac{p_z}{|\vec{p}|}}{1 - \frac{p_z}{|\vec{p}|}} \right) \\
&= \frac{1}{2} \ln \left(\frac{1 + \cos \theta}{1 - \cos \theta} \right) \\
&= \frac{1}{2} \ln \left(\frac{1 + \cos^2 \frac{\theta}{2} - \sin^2 \frac{\theta}{2}}{1 - \cos^2 \frac{\theta}{2} + \sin^2 \frac{\theta}{2}} \right) \\
&= \ln \left[\left(\frac{2 \cos^2 \frac{\theta}{2}}{2 \sin^2 \frac{\theta}{2}} \right)^{\frac{1}{2}} \right] \\
&= -\ln \tan \left(\frac{\theta}{2} \right)
\end{aligned}$$

where \vec{p} is the particle's momentum vector. The pseudo-rapidity is a good approximation for rapidity when particles are approximately massless, i.e. when $E \gg |\vec{p}|$, which almost always holds for final state particles interacting with the CMS detector. Pseudo-rapidity is a purely geometric quantity (depending only on the polar angle) and thus easy to compute. Combined with the convenient property of approximate conservation of pseudo-rapidity differences, it is a useful way of characterizing the momenta of particles going through the CMS detector.

Invariant mass

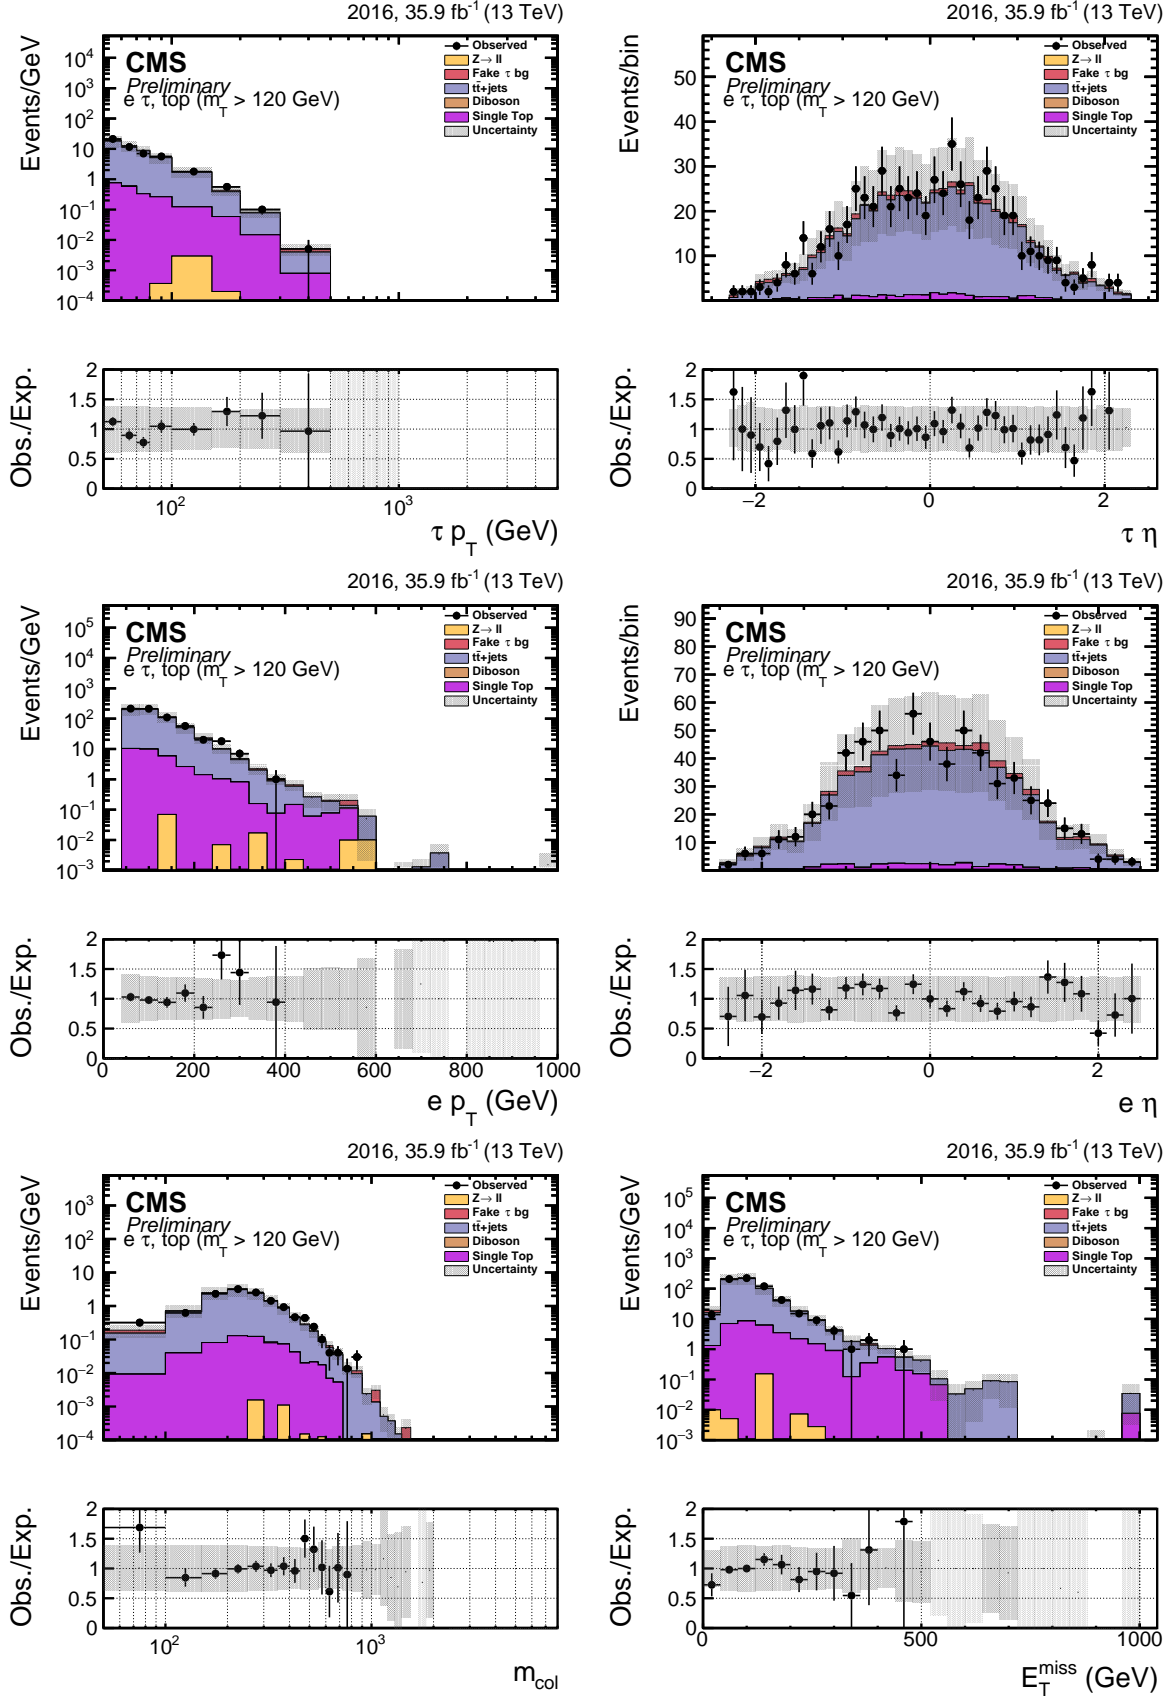
When we suspect that two final state particles observed in CMS come from the decay of the same heavy unstable particle, it is useful to reconstruct the mass of that heavy particle. This is done by calculating the **invariant mass** m_{inv} of the two final state particles. If the two particles have 4-momenta $p_1^\mu = (E_1, \vec{p}_1)$ and $p_2^\mu = (E_2, \vec{p}_2)$, then m_{inv} is calculated by summing the two 4-momenta and calculating the rest mass corresponding to the resulting 4-momentum:

$$m_{\text{inv}}^2 = (p_1^\mu + p_2^\mu)^2 = (E_1 + E_2)^2 - (\vec{p}_1 + \vec{p}_2)^2$$

This can be generalized to three or more particles.

Validation of the $t\bar{t}$ background estimation

Events with two or more b-tagged jets (the CSVv2 algorithm is used in 2016 and the DeepCSV algorithm in 2017-18) are enriched in $t\bar{t}$ background. The figures in this appendix compare data distributions with expectations in two distinct regions ($m_T < 120$ GeV and $m_T > 120$ GeV) for each tau channel, and for all years 2016-18. They all show good agreement, within uncertainties, validating the $t\bar{t}$ background estimation.



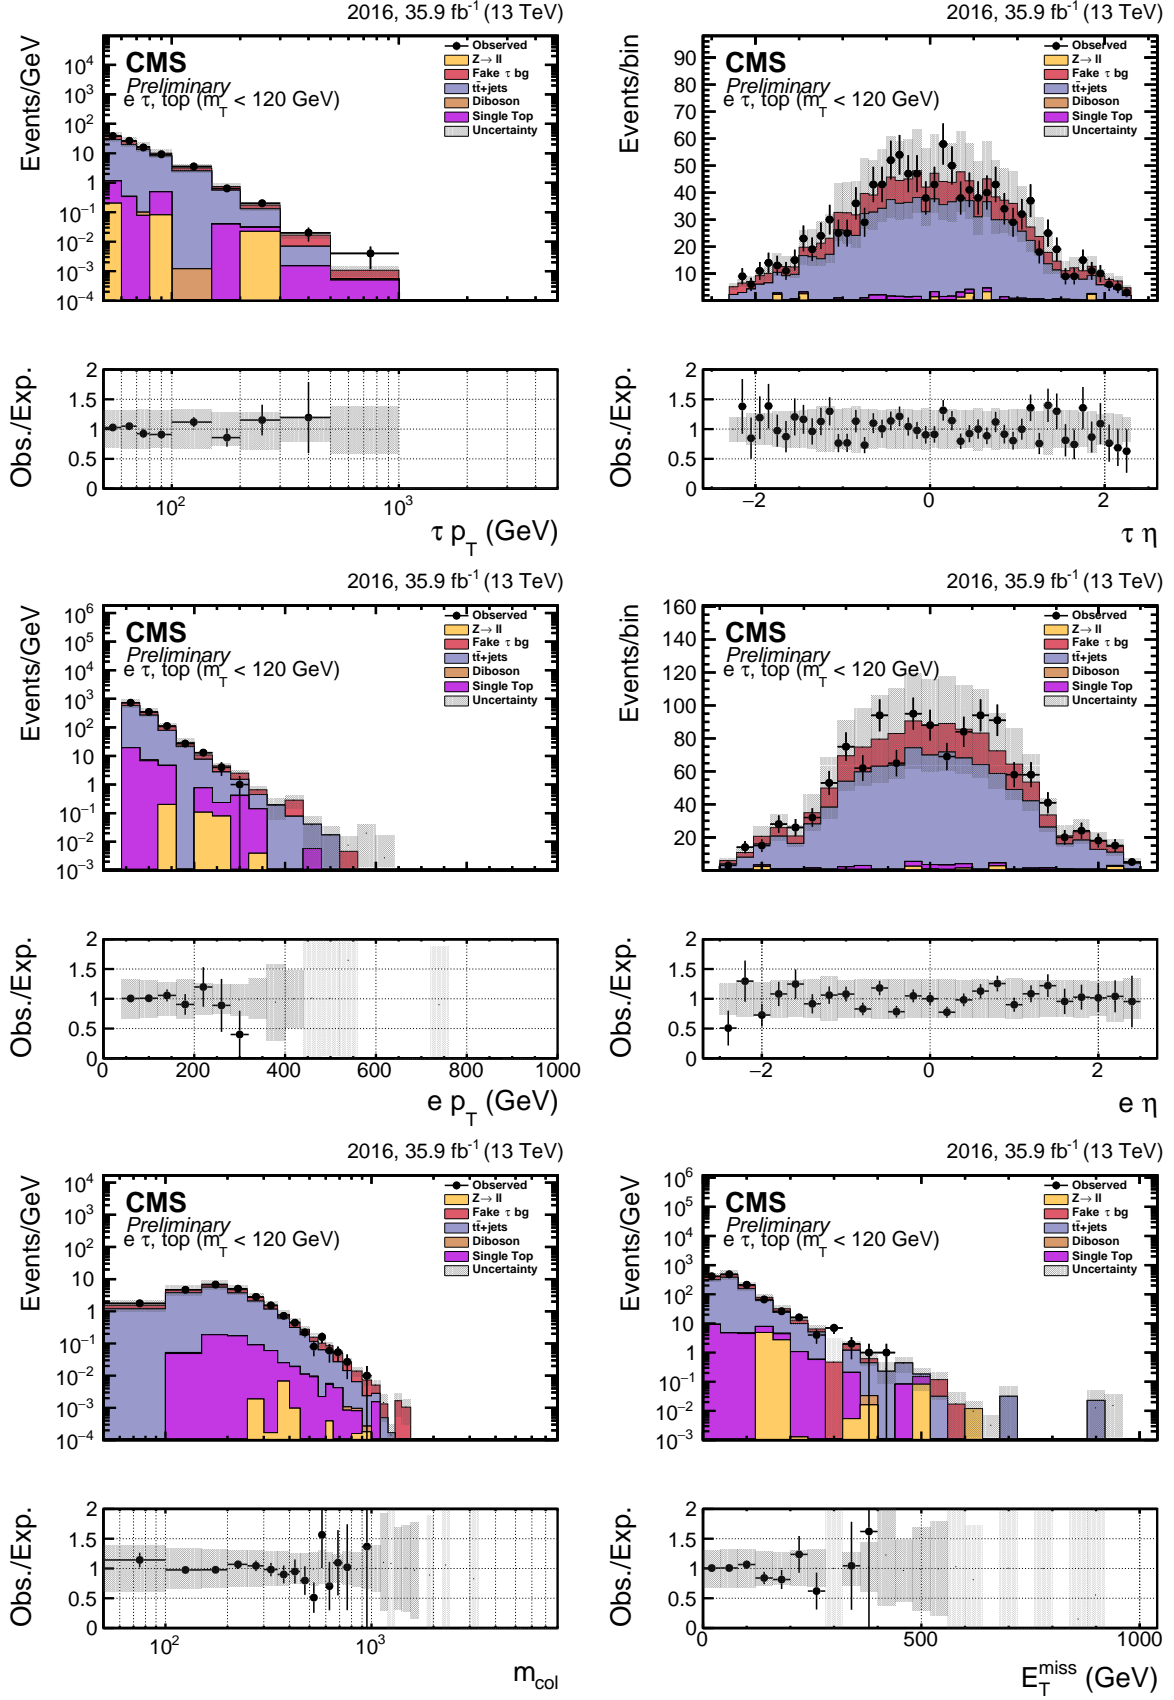


Figure – $t\bar{t}$ ($m_T < 120$ GeV) control region distributions, $e\tau$ channel, 2016. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(e)$, (4) $\eta(e)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

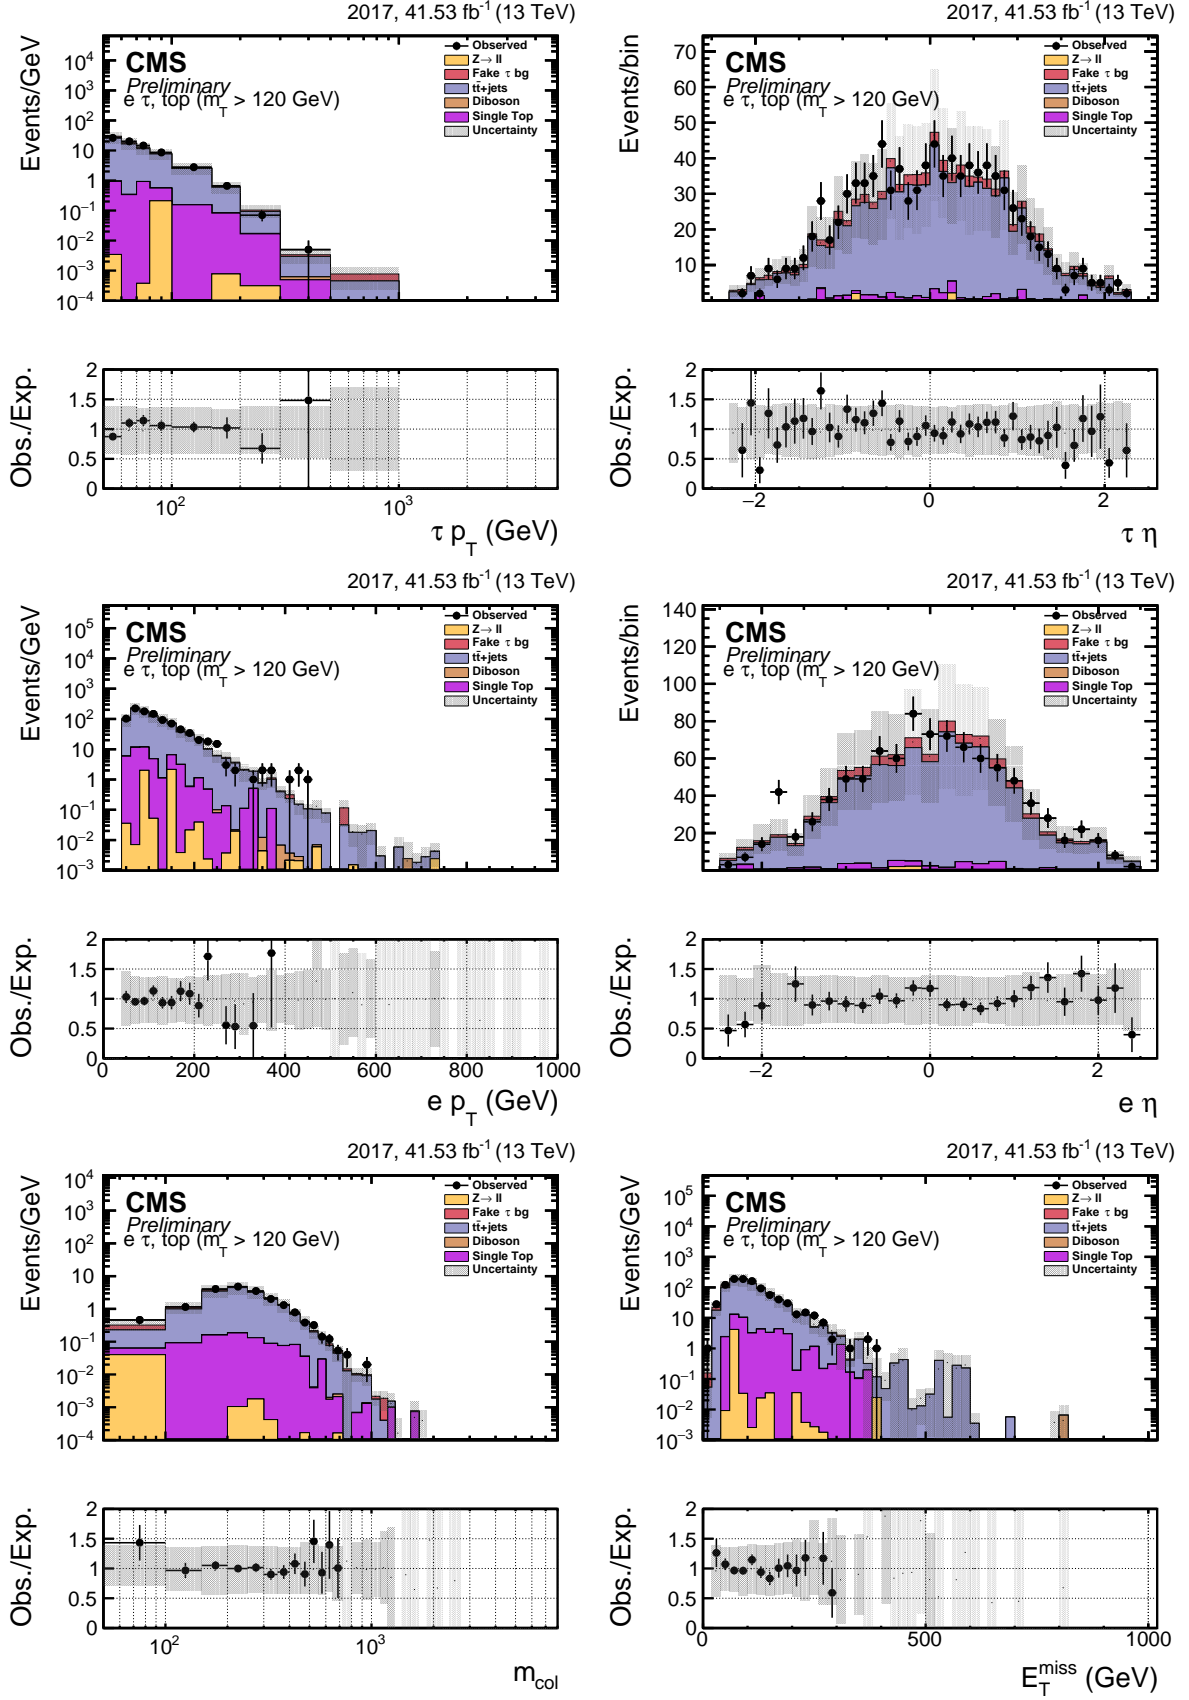


Figure – $t\bar{t}$ ($m_T > 120$ GeV) control region distributions, $e\tau$ channel, 2017. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(e)$, (4) $\eta(e)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

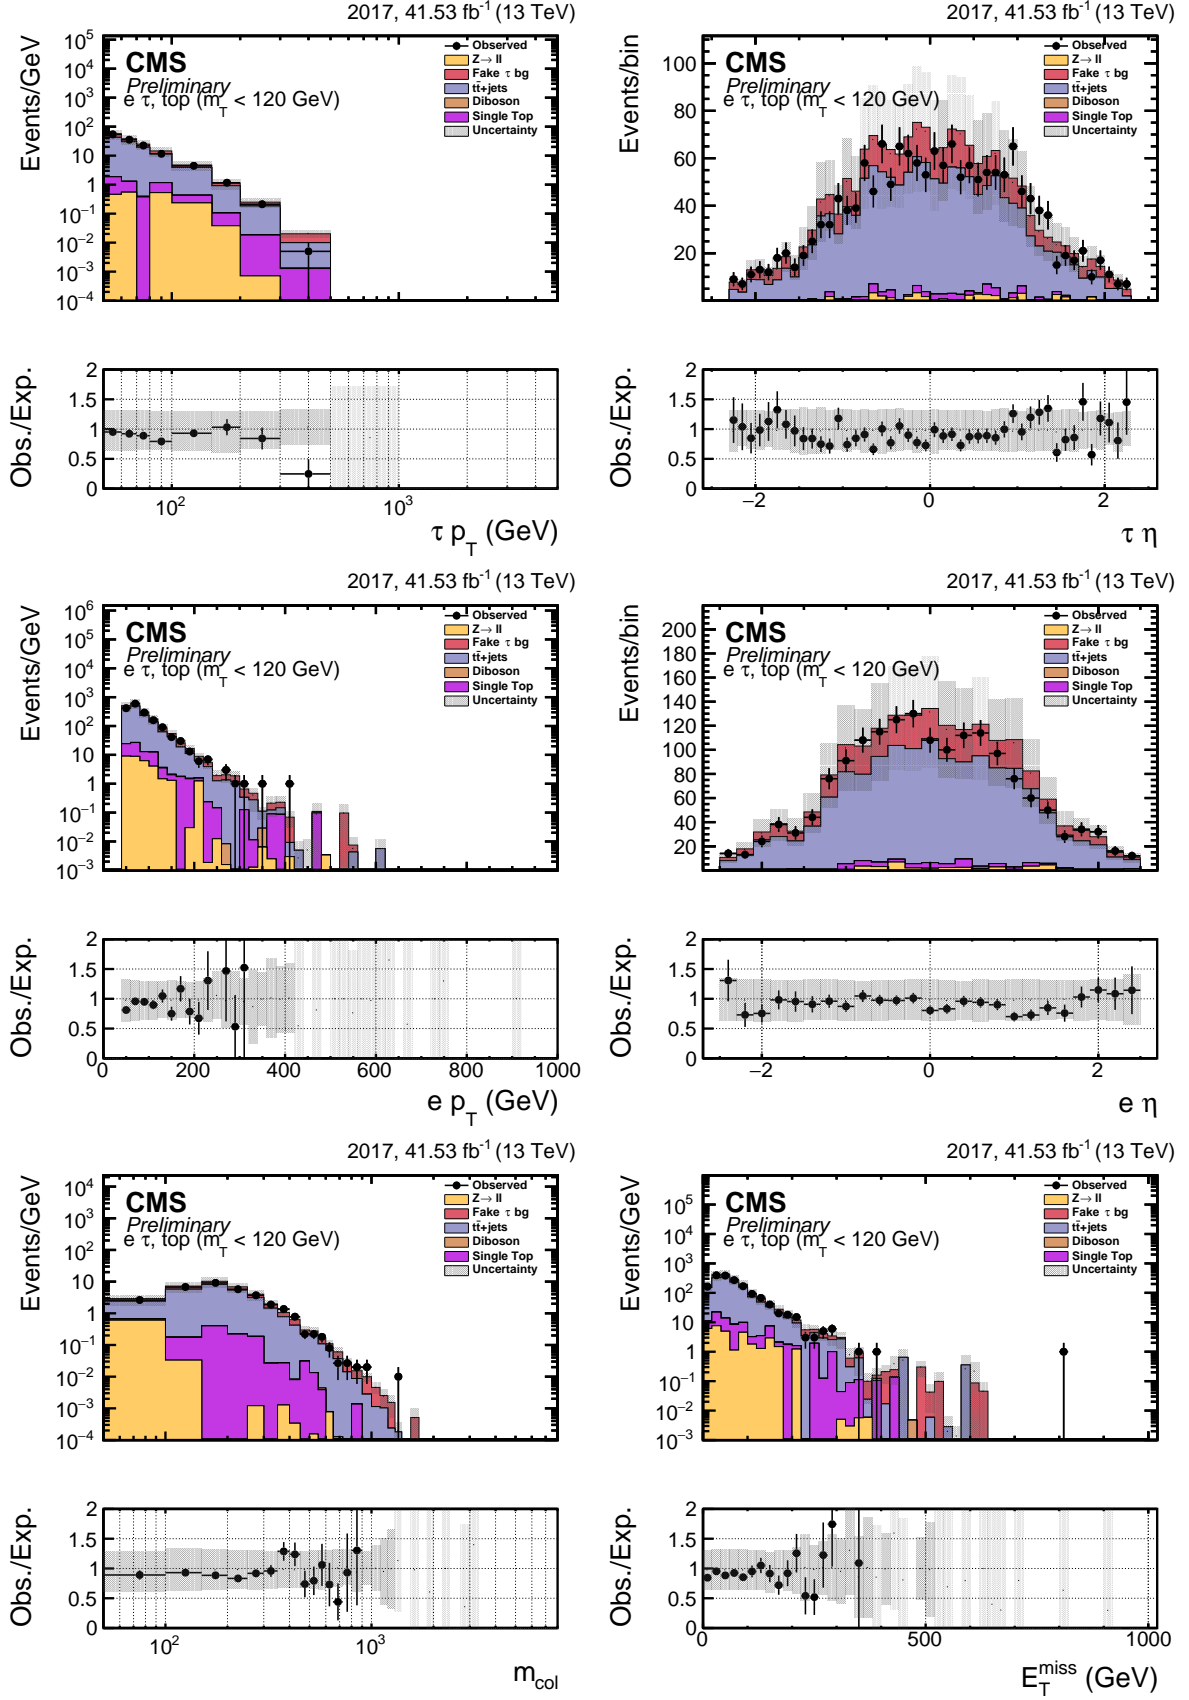


Figure – $t\bar{t}$ ($m_T < 120$ GeV) control region distributions, $e\tau$ channel, 2017. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(e)$, (4) $\eta(e)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

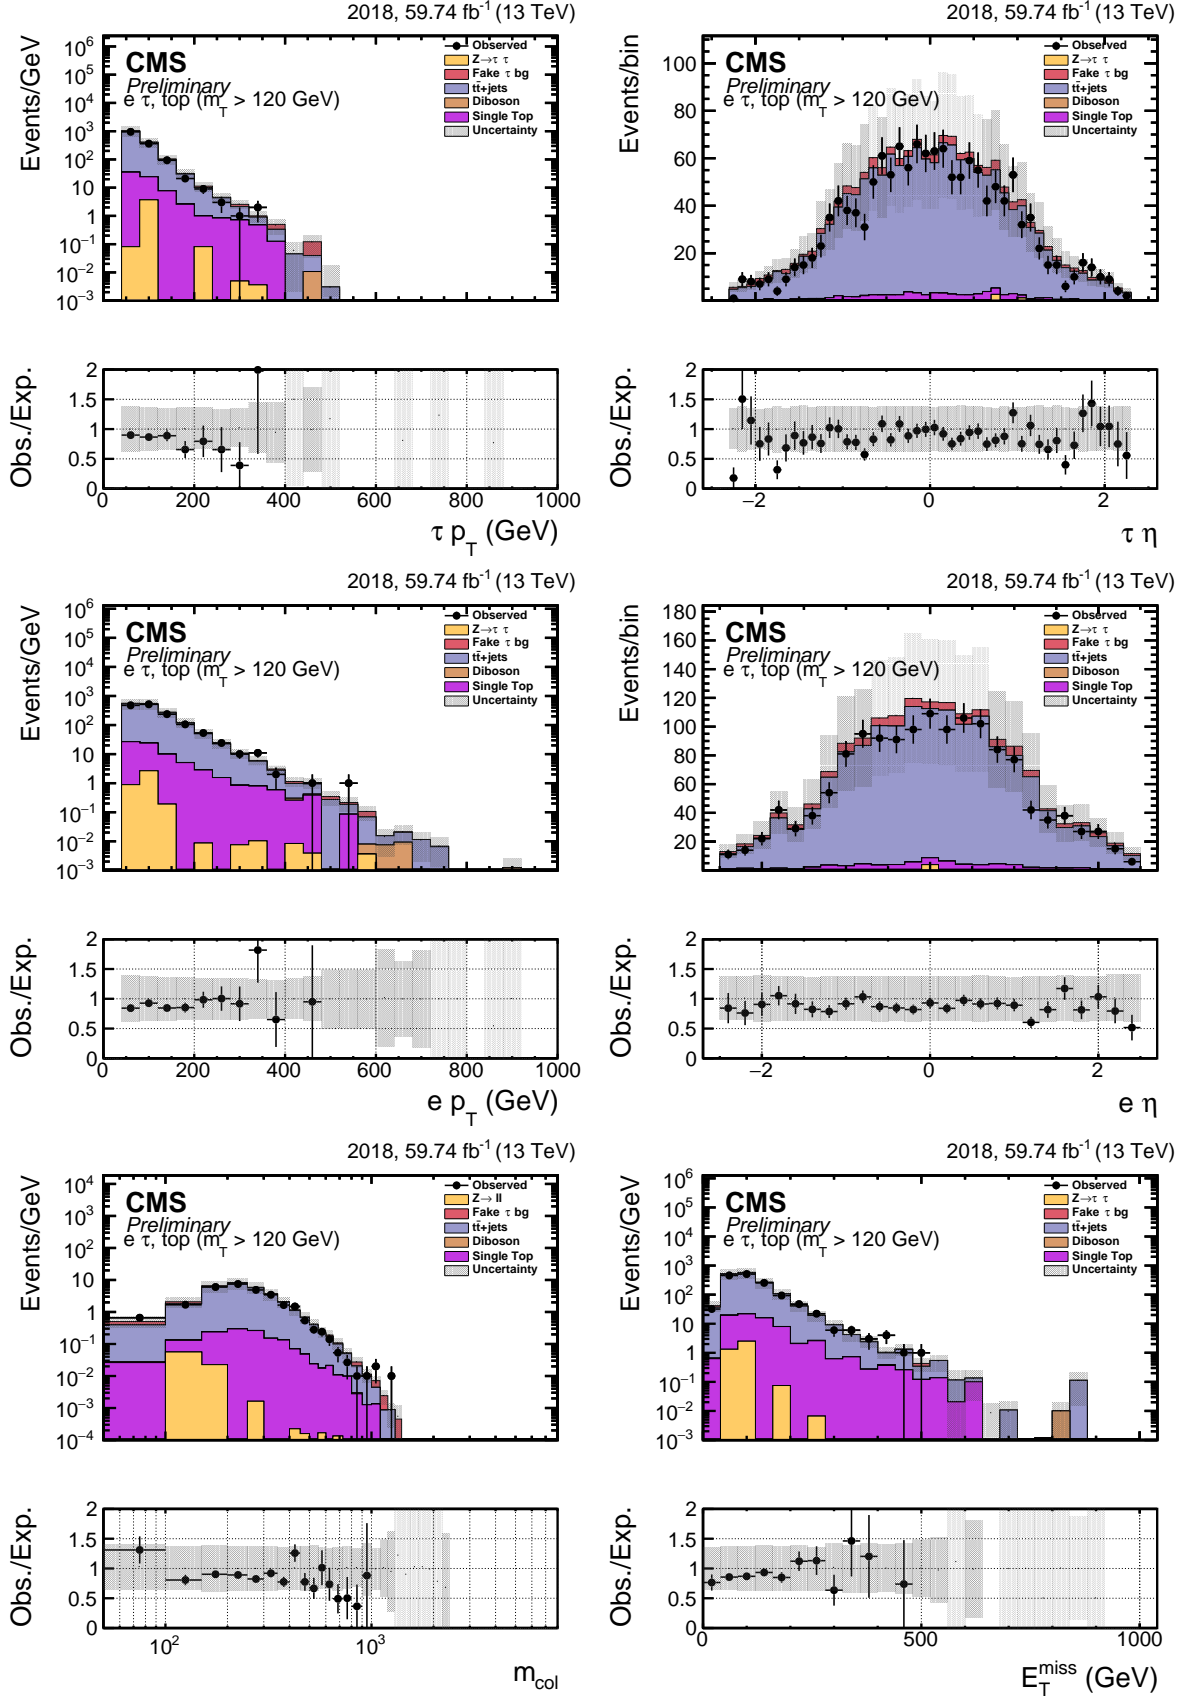


Figure – $t\bar{t}$ ($m_T > 120$ GeV) control region distributions, $e\tau$ channel, 2018. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(e)$, (4) $\eta(e)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

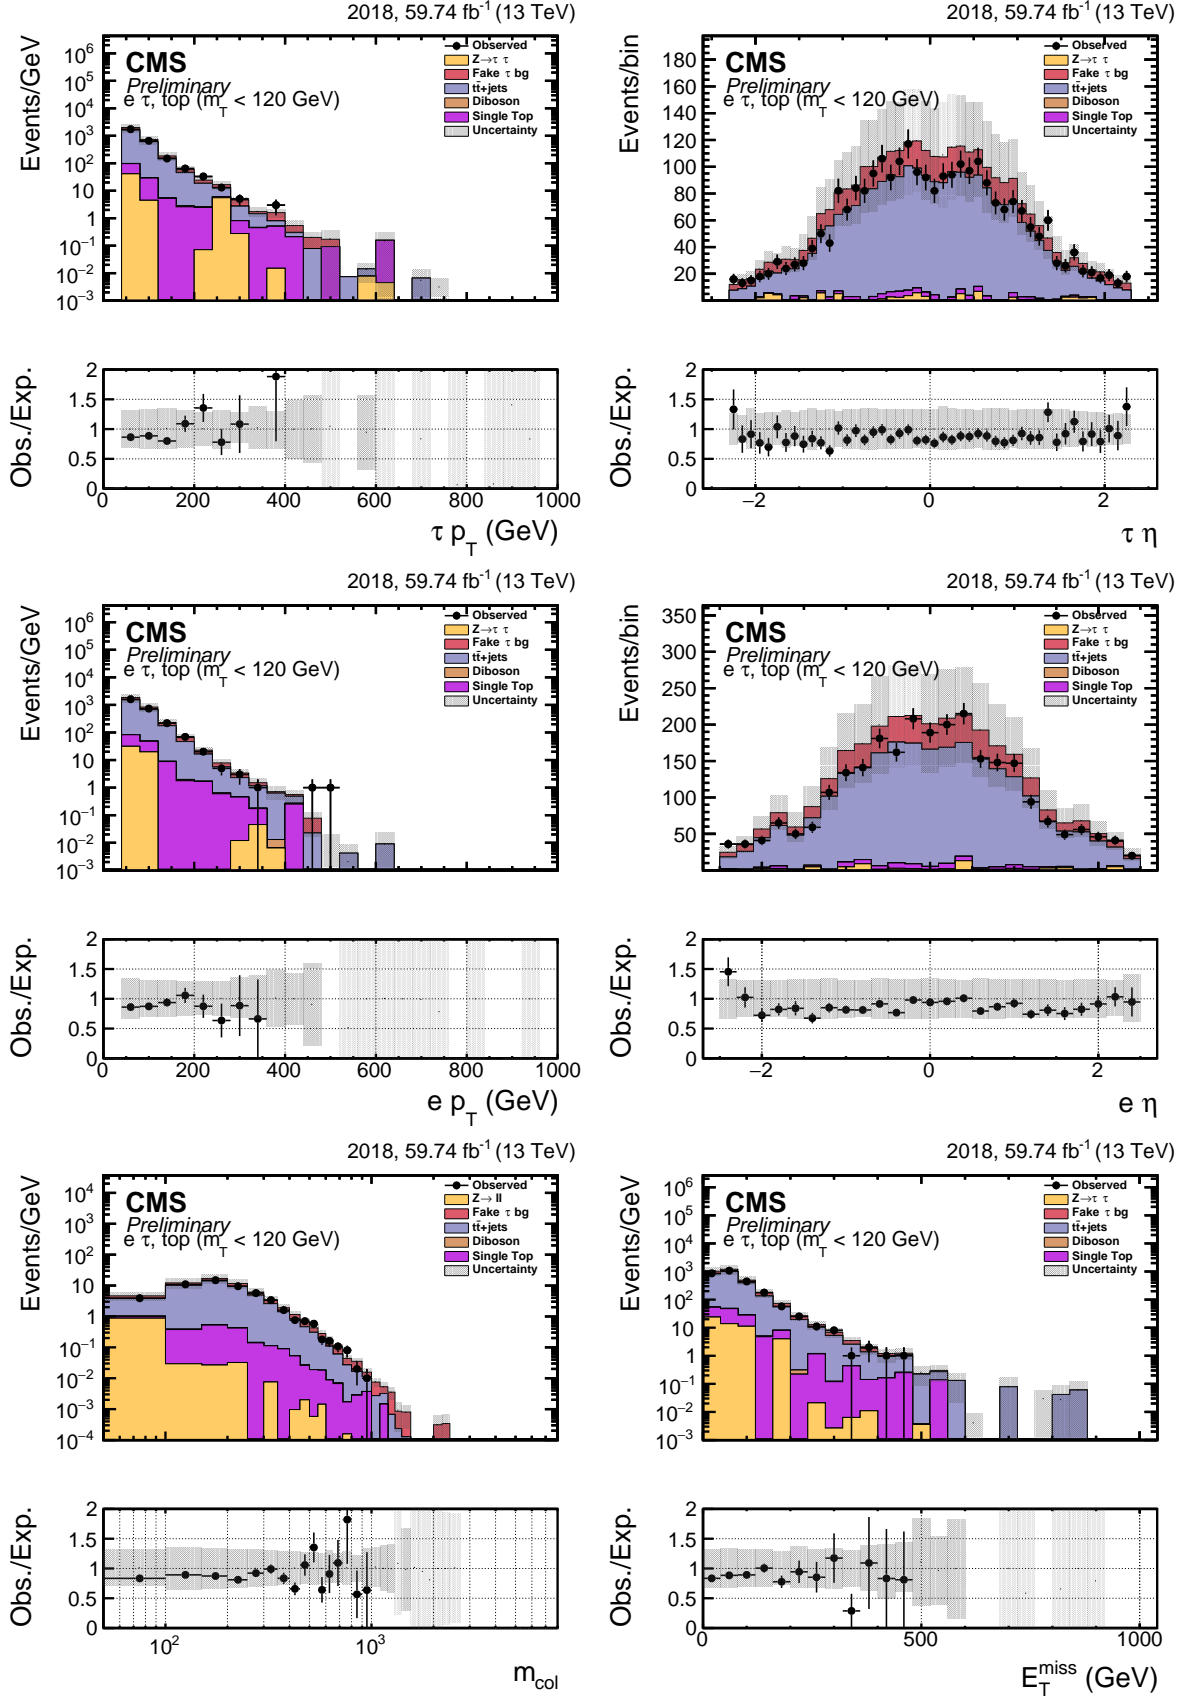


Figure – $t\bar{t}$ ($m_T < 120$ GeV) control region distributions, $e\tau$ channel, 2018. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(e)$, (4) $\eta(e)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

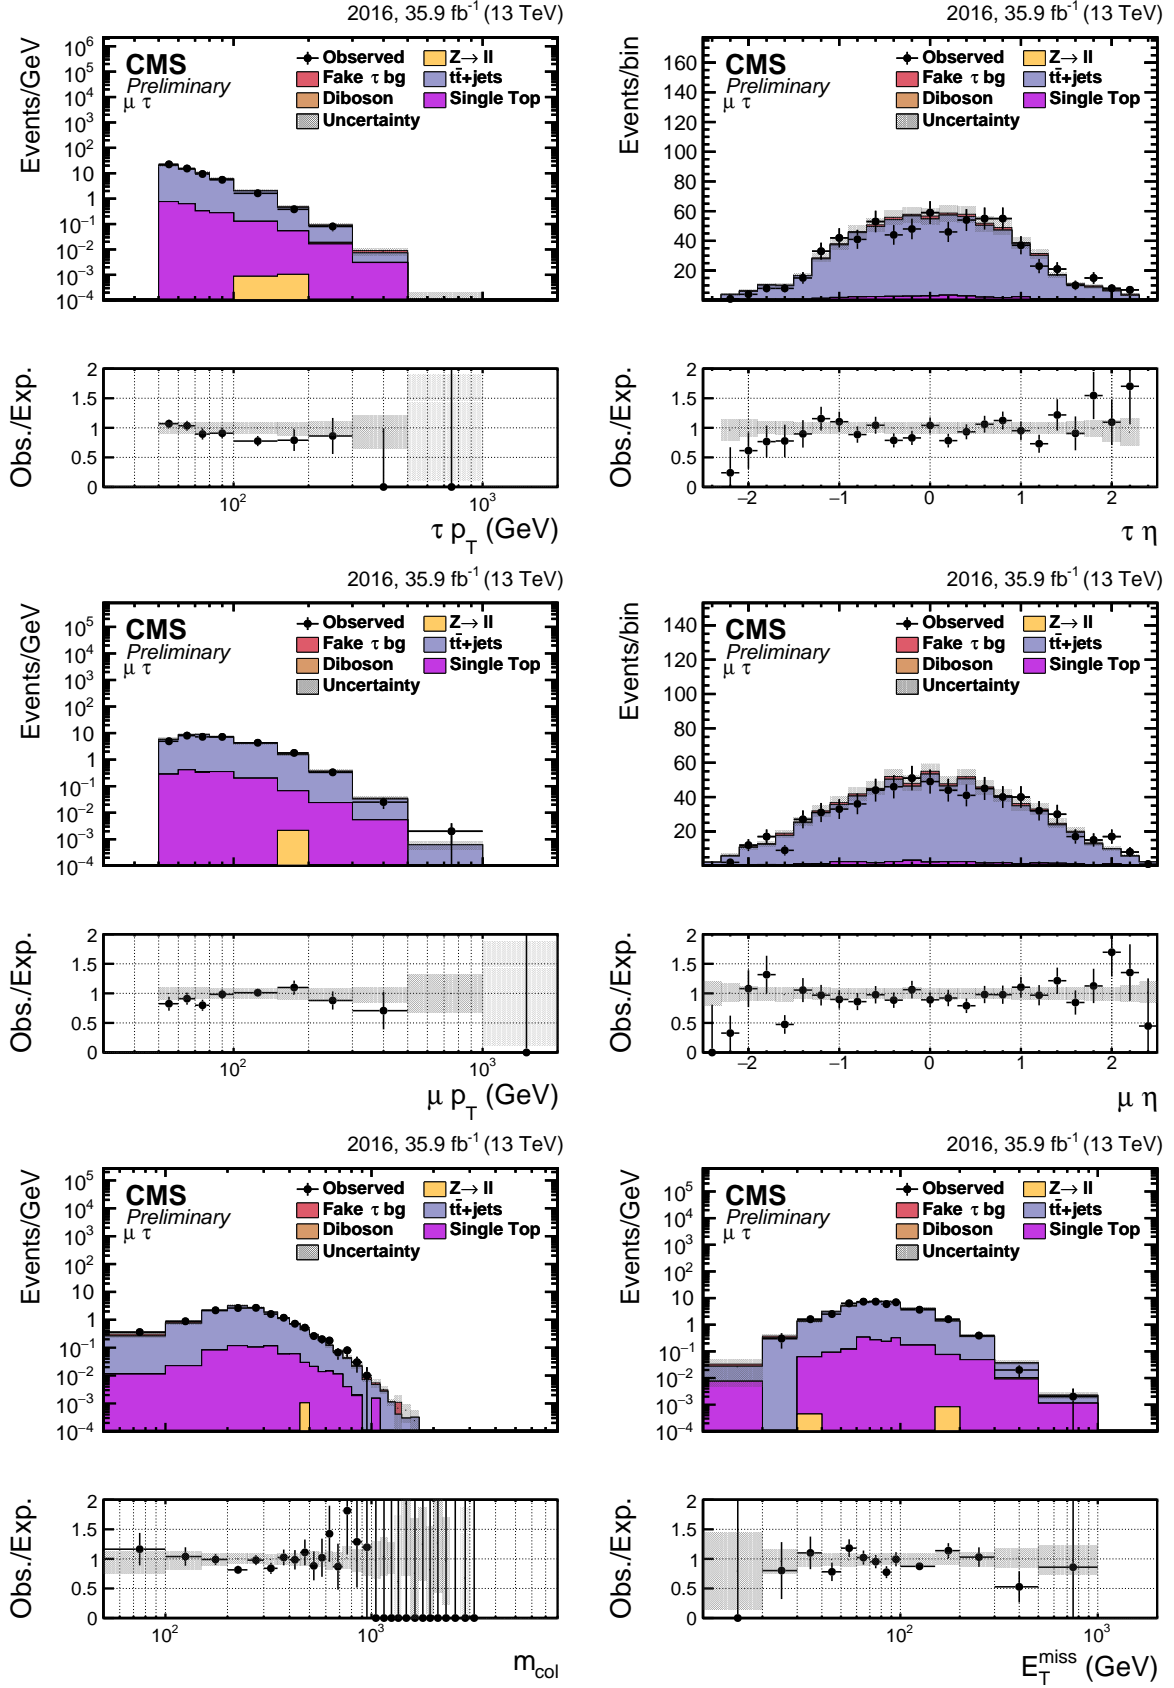


Figure – $t\bar{t}$ ($m_T > 120$ GeV) control region distributions, $\mu\tau$ channel, 2016. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

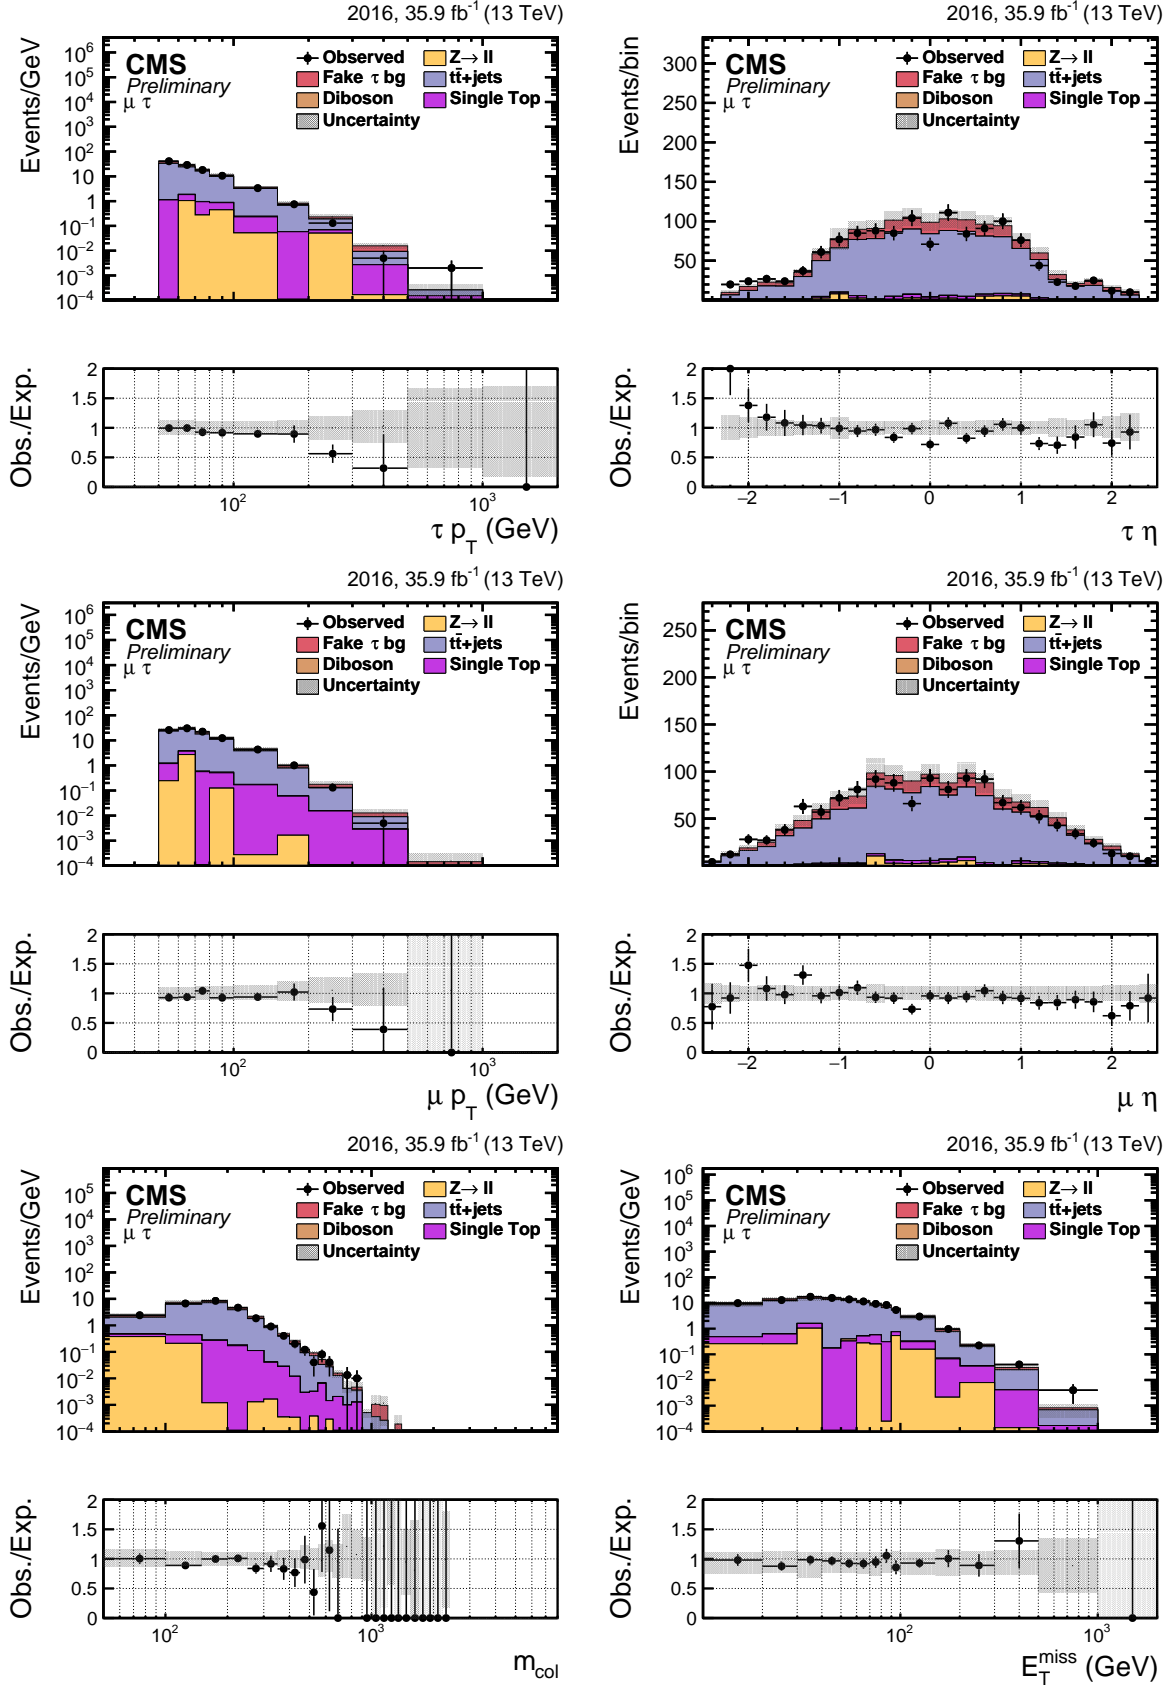


Figure – $t\bar{t}$ ($m_T < 120$ GeV) control region distributions, $\mu\tau$ channel, 2016. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

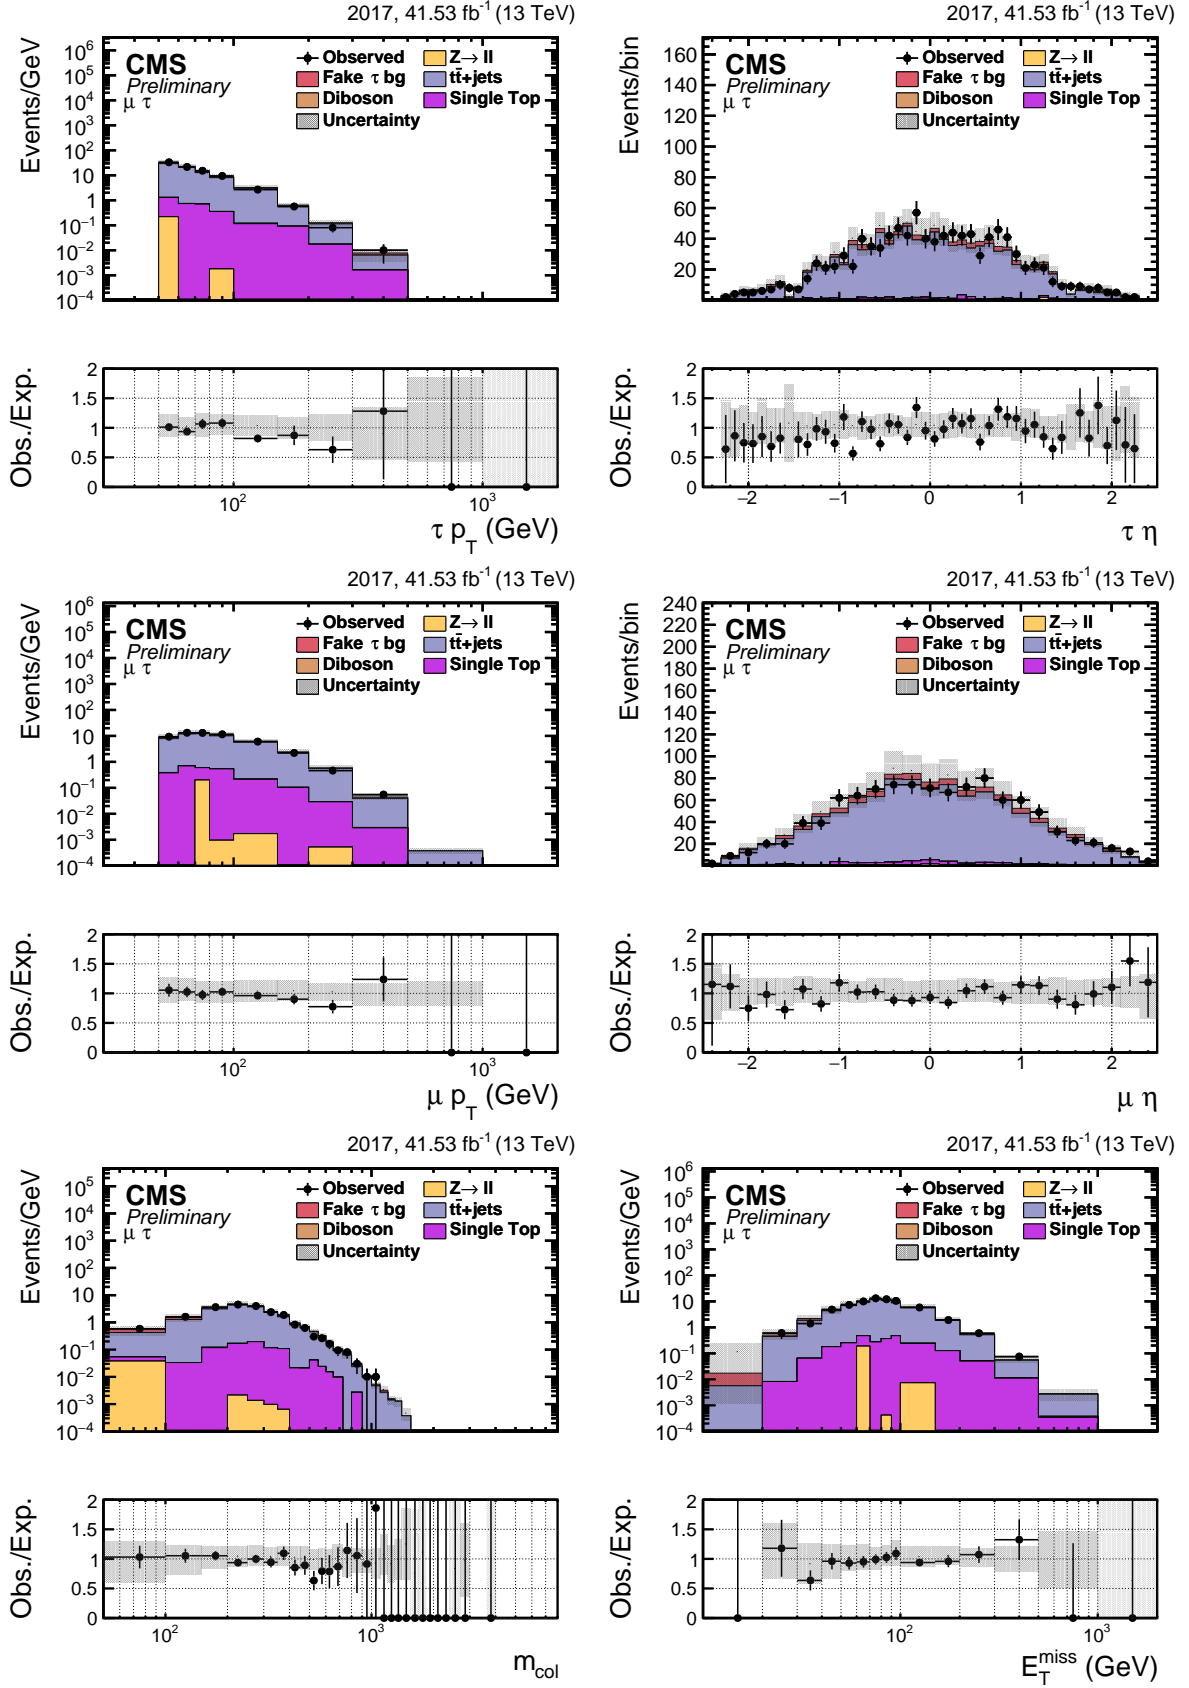


Figure – $t\bar{t}$ ($m_T > 120$ GeV) control region distributions, $\mu\tau$ channel, 2017. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

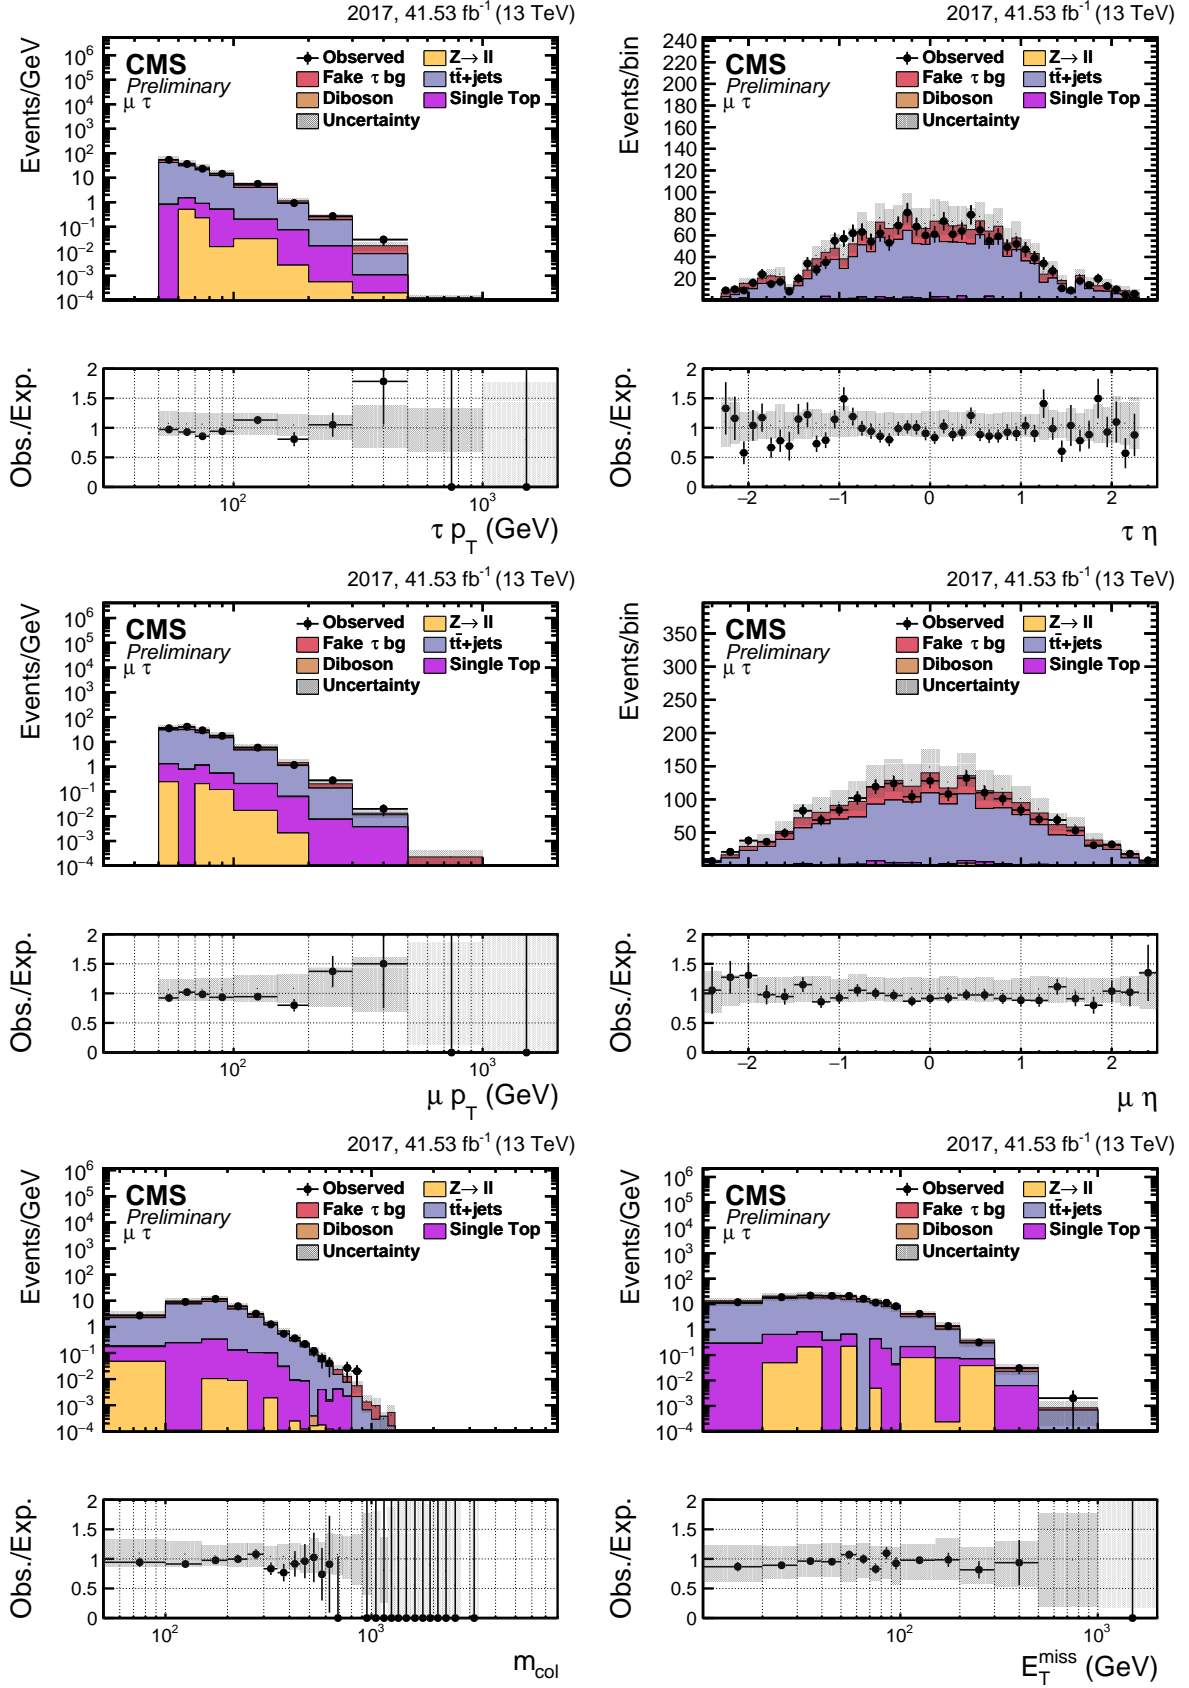


Figure – $t\bar{t}$ ($m_T < 120$ GeV) control region distributions, $\mu\tau$ channel, 2017. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

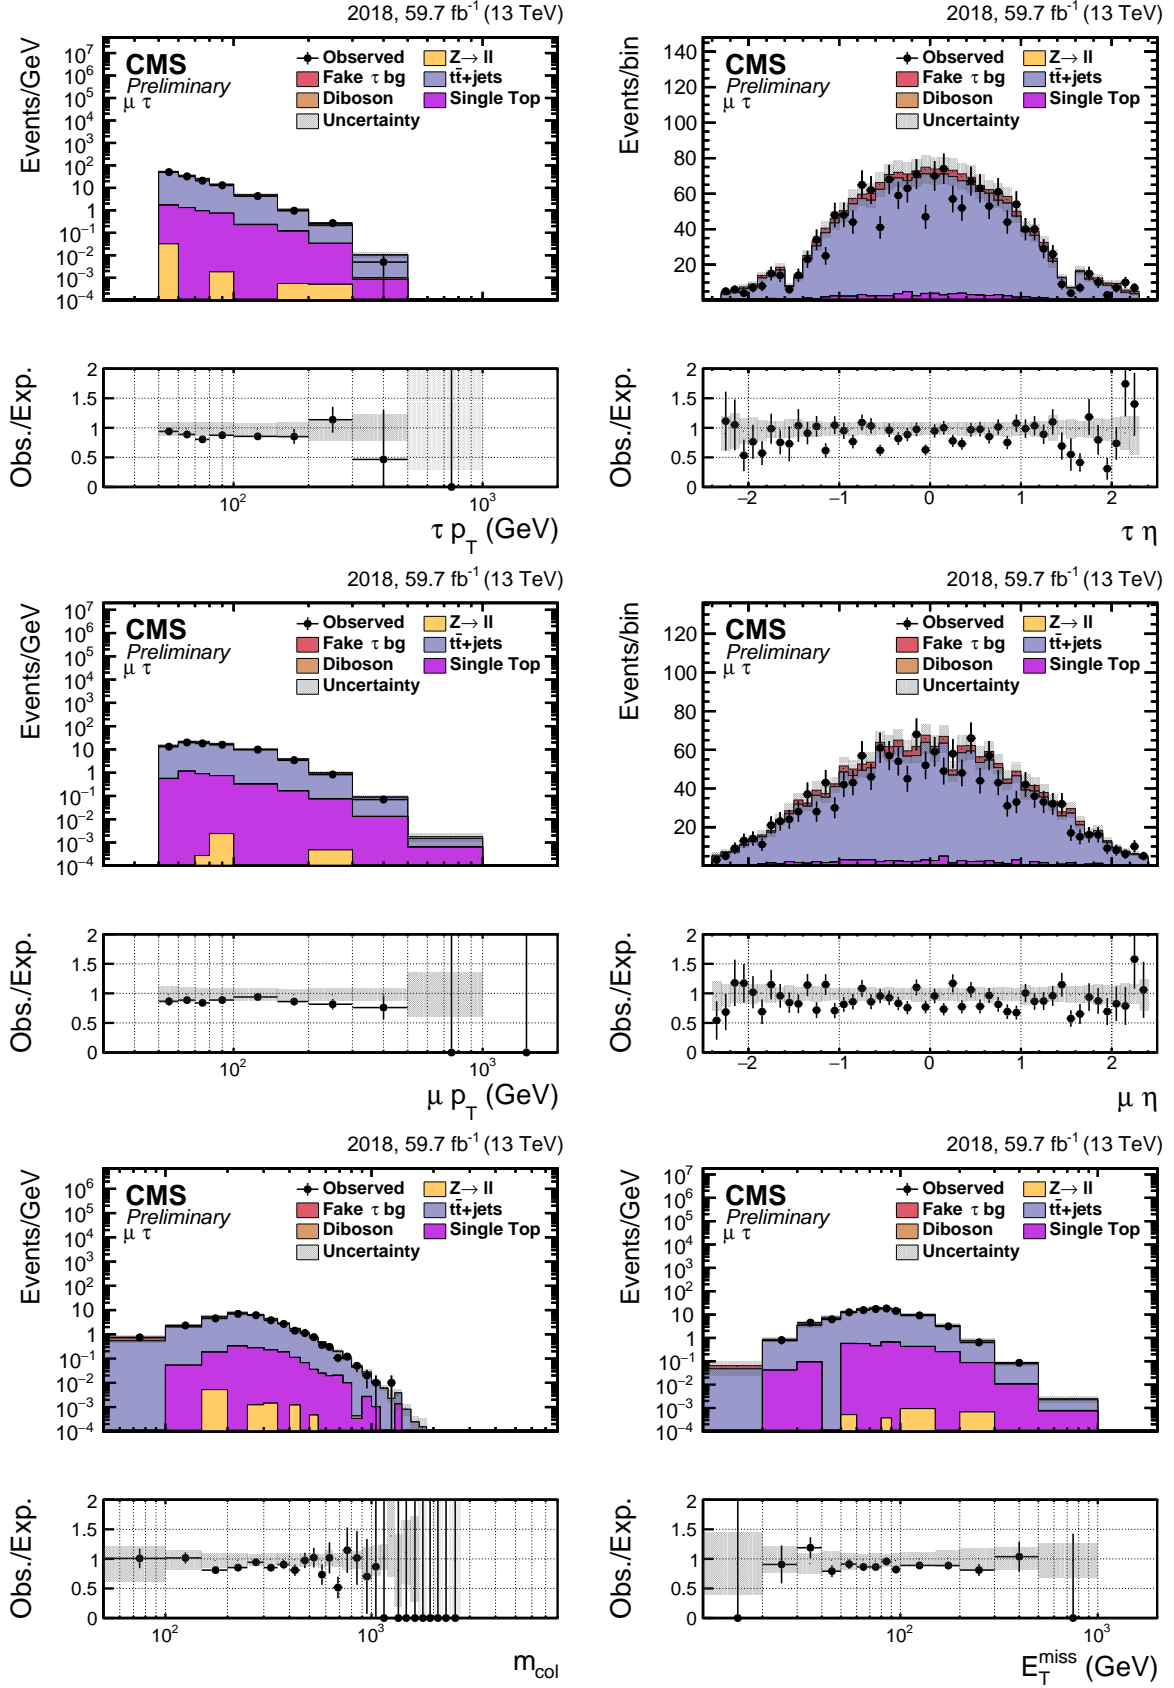


Figure – $t\bar{t}$ ($m_T > 120$ GeV) control region distributions, $\mu\tau$ channel, 2018. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

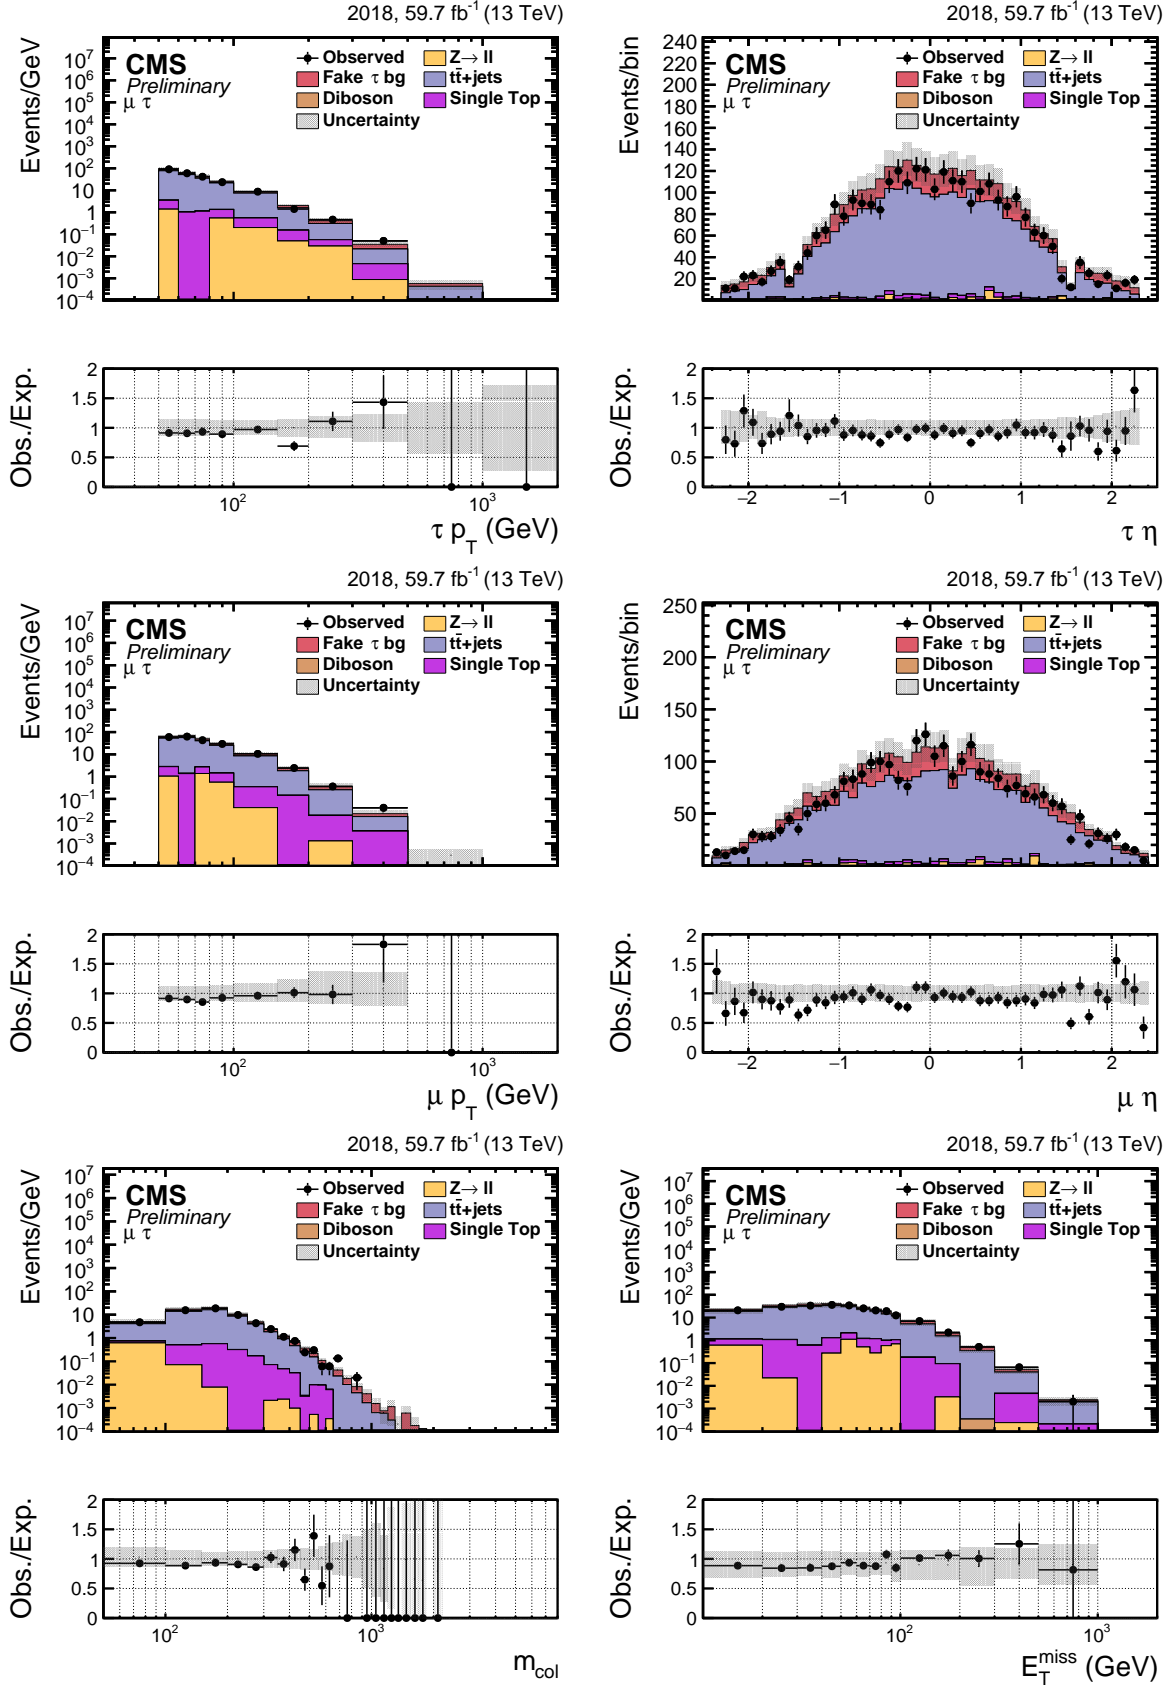


Figure – $t\bar{t}$ ($m_T < 120$ GeV) control region distributions, $\mu\tau$ channel, 2018. From left to right, top to bottom: (1) $p_T(\tau_h)$, (2) $\eta(\tau_h)$, (3) $p_T(\mu)$, (4) $\eta(\mu)$ (5) collinear mass, and (6) MET. The shaded uncertainty band represents both statistical and systematic uncertainties.

Comparison of jet $\rightarrow \tau$ fakes in the $e\tau$ and $\mu\tau$ channels

It is evident when looking at the mass plots in figure 9.6 that there are more jet $\rightarrow \tau$ fakes in the $e\tau$ channel than in the $\mu\tau$ channel. Since the tau selections are the same in both channels, the difference must come from the selection of the light lepton (electron or muon). Muons are very reliably reconstructed by CMS, while electrons are more difficult to reconstruct and identify due to their high interaction rates with the silicon tracker. We thus made the hypothesis that the extra events in the $e\tau$ channel are “double fakes” i.e. events in which both the reconstructed electron and tau are in fact misidentified jets.

This hypothesis is supported by the collinear mass plots shown in the figure at the end of this appendix. m_{col} distributions are shown in the $\mu\tau$ and $e\tau$ channels using W+jets simulation samples instead of the data-driven jet $\rightarrow \tau$ background. In the $\mu\tau$ channel, the W+jets events are enough to provide a good agreement with the data distribution, while in the $e\tau$ channel some events are clearly missing. W+jets events passing tau channel selections are mostly “single fake” with the W boson decaying to an electron or muon and a jet faking a tau. Other possibilities¹ are much rarer. “Double fakes” events are more likely to come from the QCD multijets process. They can be estimated separately from W+jets events with a “same sign” data-driven method, similar to the one employed to cross-check the fake rates in the $e\mu$ channel (see chapter 8). A subset of the $e\tau$ signal region is defined, by requiring the electron and the tau to have the same signs. Simulated same sign events with real electrons and taus are subtracted from the data distribution, in order to leave only “double fakes” events. The yield thus obtained is multiplied by two, in order to obtain an estimate for the whole signal region, which has no sign selection. The resulting distribution is added to the other backgrounds and compared to the data distribution. The agreement is good, as shown in the figure.

1. For example “double fakes” events or events with the W decaying to a tau and a jet faking an electron or a muon.

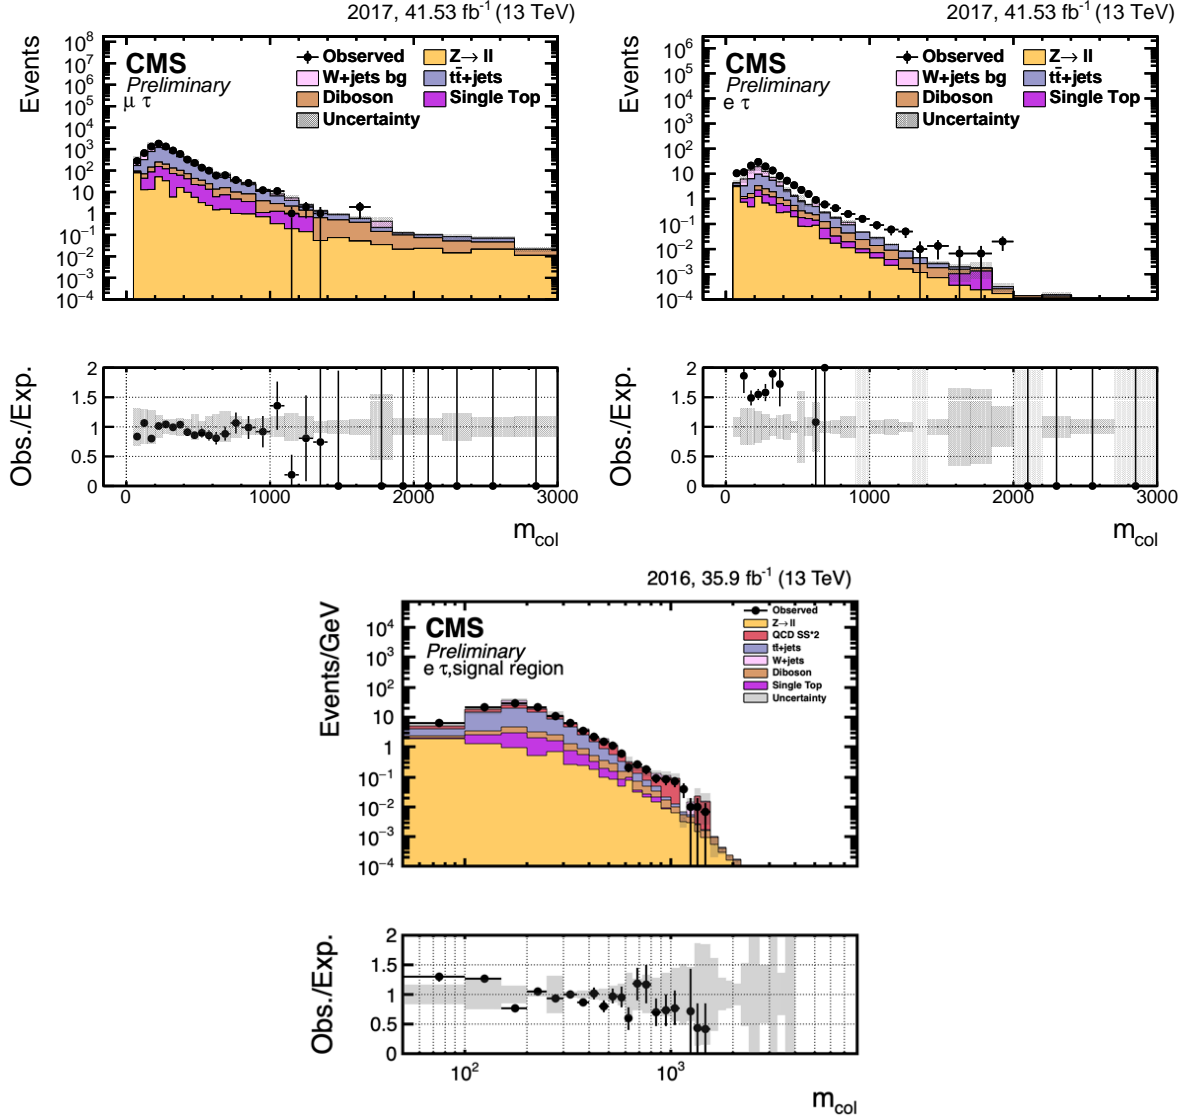


Figure – On the top left, the m_{col} distribution is shown in the $\mu\tau$ channel, with the W+jets simulation replacing the data-driven jet $\rightarrow \tau$ fakes estimation. On the top, the same is shown for the $e\tau$ channel. The agreement between data and expectations is much better in the $\mu\tau$ channel, showing that a significant number of events are missing in the $e\tau$ channel. The bottom plot adds a data-driven estimation of the QCD “double fakes” background to the $e\tau$ channel, providing a much better agreement with data.

Bibliography

- [1] J. J. Thomson, “Cathode Rays”, *Philosophical Magazine* (1897) 293.
- [2] M. Gell-Mann, “The interpretation of the new particles as displaced charge multiplets”, *Il Nuovo Cimento* **4** (1956) 848.
- [3] T. Nakano and K. Nishijima, “Charge Independence for V-particles”, *Progress of Theoretical Physics* **10** (1953) 581.
- [4] J. Gilman, “High Energy Electroproduction”, in *Proceedings of the Lepton-Photon Symposium*, SLAC, pp. 1–55, 1969.
- [5] H. Fritzsch, M. Gell-Mann and H. Leutwyler, “Advantages of the color octet gluon picture”, *Physics Letters B* **47** (1973) 365.
- [6] R. P. Feynman, “Mathematical Formulation of the Quantum Theory of Electromagnetic Interaction”, *Physical Review* **80** (1950) 440.
- [7] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons”, *Physical Review Letters* **13** (1964) 321.
- [8] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons”, *Physical Review Letters* **13** (1964) 508.
- [9] S. Weinberg, “A Model of Leptons”, *Physical Review Letters* **19** (1967) 1264.
- [10] Super-Kamiokande Collaboration, “Evidence for oscillation of atmospheric neutrinos”, *Physical Review Letters* **81** (1998) 1562 [[hep-ex/9807003](#)].
- [11] M. Kobayashi and T. Maskawa, “CP-Violation in the Renormalizable Theory of Weak Interaction”, *Progress of Theoretical Physics* **49** (1973) 652.
- [12] Y. Kuno and Y. Okada, “Muon Decay and Physics Beyond the Standard Model”, *Reviews of Modern Physics* **73** (2001) 151.
- [13] KATRIN Collaboration, “An improved upper limit on the neutrino mass from a direct kinematic method by KATRIN”, *Physical Review Letters* **123** (2019) 221802.
- [14] S. A. Thomas, F. B. Abdalla and O. Lahav, “Upper Bound of 0.28 eV on the Neutrino Masses from the Largest Photometric Redshift Survey”, *Physical Review Letters* **105** (2010) 031301.
- [15] Planck Collaboration, “Planck 2013 results. XVI. Cosmological parameters”, *Astronomy & Astrophysics* **571** (2014) A16.
- [16] TASSO Collaboration, “TASSO results on e^+e^- annihilation between 13 and 31.6 GeV and evidence for three jet events”, in *Proceedings of the Lepton-Photon Symposium*, Chicago, pp. 34–51, 1979.

- [17] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”, *Physics Letters B* **716** (2012) 30.
- [18] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Physics Letters B* **716** (2012) 1.
- [19] A. Shomer, “A pedagogical explanation for the non-renormalizability of gravity”, arXiv note, 2007 [[arXiv:0709.3555](https://arxiv.org/abs/0709.3555)], (<http://arxiv.org/abs/0709.3555>).
- [20] J. E. Kim, “Light pseudoscalars, particle physics and cosmology”, *Physics Reports* **150** (1987) 1.
- [21] D. Dubbers and M. G. Schmidt, “The neutron and its role in cosmology and particle physics”, *Reviews of Modern Physics* **83** (2011) 1111.
- [22] A. G. Cohen, A. De Rujula and S. L. Glashow, “A Matter-Antimatter Universe?”, *The Astrophysical Journal* **495** (1998) 539.
- [23] M. Trodden, “Electroweak Baryogenesis: A Brief Review”, in *Proceedings of the XXXIIIrd Rencontres de Moriond*, 1998, [[arXiv:hep-ph/9805252](https://arxiv.org/abs/hep-ph/9805252)].
- [24] F. Zwicky, “On the Masses of Nebulae and of Clusters of Nebulae”, *The Astrophysical Journal* **86** (1937) 217.
- [25] K. Freese, “Review of Observational Evidence for Dark Matter in the Universe and in upcoming searches for Dark Stars”, in *EAS Publications Series*, vol. 36, Lyon, pp. 113–126, 2008, [[arXiv:0812.4005](https://arxiv.org/abs/0812.4005)].
- [26] K. Blum, “Dark Matter: Theory Perspective”, presentation at EPS-HEP Ghent, 2019.
- [27] B. Carr, F. Kuhnel and M. Sandstad, “Primordial Black Holes as Dark Matter”, *Physical Review D* **94** (2016) 083504.
- [28] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiattia, A. Diercks, P. M. Garnavich et al., “Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant”, *The Astronomical Journal* **116** (1998) 1009.
- [29] I. L. Shapiro and J. Sola, “Cosmological Constant Problems and Renormalization Group”, *Journal of Physics A: Mathematical and Theoretical* **40** (2007) 6583.
- [30] L. Susskind, “Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory”, *Physical Review D* **20** (1979) 2619.
- [31] M. Shaposhnikov and C. Wetterich, “Asymptotic safety of gravity and the Higgs boson mass”, *Physics Letters B* **683** (2010) 196.
- [32] S. P. Martin, “A Supersymmetry Primer”, arXiv note, 1998 [[arXiv:hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356)].
- [33] R. Barbier, C. Berat, M. Besancon, M. Chemtob, A. Deandrea, E. Dudas et al., “R-parity violating supersymmetry”, *Physics Reports* **420** (2005) 1.
- [34] F. D. Steffen, “Supersymmetric Dark Matter Candidates - The Lightest Neutralino, the Gravitino, and the Axino”, arXiv note, 2007 [[arXiv:0711.1240](https://arxiv.org/abs/0711.1240)].
- [35] G. C. Branco, P. M. Ferreira, L. Lavoura, M. N. Rebelo, M. Sher and J. P. Silva, “Theory and phenomenology of two-Higgs-doublet models”, *Physics Reports* **516** (2012) 1.
- [36] N. Arkani-Hamed, S. Dimopoulos and G. Dvali, “The Hierarchy Problem and New Dimensions at a Millimeter”, *Physics Letters B* **429** (1998) 263.
- [37] L. Randall and R. Sundrum, “A Large Mass Hierarchy from a Small Extra Dimension”, *Physical Review Letters* **83** (1999) 3370.
- [38] D. M. Eardley and S. B. Giddings, “Classical Black Hole Production in High-Energy Collisions”, *Physical Review D* **66** (2002) 044011.

- [39] S. D. H. Hsu, “Quantum Production of Black Holes”, *Physics Letters B* **555** (2003) 92.
- [40] S. W. Hawking, “Particle creation by black holes”, *Communications in Mathematical Physics* **43** (1975) 199.
- [41] S. Dimopoulos and G. Landsberg, “Black Holes at the LHC”, *Physical Review Letters* **87** (2001) 161602.
- [42] L. A. Anchordoqui, J. L. Feng, H. Goldberg and A. D. Shapere, “Inelastic Black Hole Production and Large Extra Dimensions”, *Physics Letters B* **594** (2004) 363.
- [43] P. Meade and L. Randall, “Black Holes and Quantum Gravity at the LHC”, *Journal of High Energy Physics* **2008** (2008) 003.
- [44] X. Calmet, W. Gong and S. D. H. Hsu, “Colorful quantum black holes at the LHC”, *Physics Letters B* **668** (2008) 20.
- [45] A. Djouadi, “The Anatomy of Electro-Weak Symmetry Breaking. II: The Higgs bosons in the Minimal Supersymmetric Model”, *Physics Reports* **459** (2008) 1.
- [46] A. Antaramian, L. J. Hall and A. Rain, “Flavor changing interactions mediated by scalars at the weak scale”, *Physical Review Letters* **69** (1992) 1871.
- [47] G. C. Branco, W. Grimus and L. Lavoura, “Relating the scalar flavour-changing neutral couplings to the CKM matrix”, *Physics Letters B* **380** (1996) 119.
- [48] R. Casalbuoni, S. De Curtis, D. Dominici and R. Gatto, “SM Kaluza-Klein Excitations and Electroweak Precision Tests”, *Physics Letters B* **462** (1999) 48.
- [49] A. Delgado, A. Pomarol and M. Quiros, “Electroweak and Flavor Physics in Extensions of the Standard Model with Large Extra Dimensions”, *Journal of High Energy Physics* **2000** (2000) 030.
- [50] H. Georgi and S. L. Glashow, “Unity of All Elementary-Particle Forces”, *Physical Review Letters* **32** (1974) 438.
- [51] H. Fritzsch and P. Minkowski, “Unified interactions of leptons and hadrons”, *Annals of Physics* **93** (1975) 193.
- [52] J. L. Hewett and T. G. Rizzo, “Low-energy phenomenology of superstring-inspired E6 models”, *Physics Reports* **183** (1989) 193.
- [53] P. Langacker, “Grand unified theories and proton decay”, *Physics Reports* **72** (1981) 185.
- [54] M. Miura, “Search for Proton Decay via $p \rightarrow e^+ \pi^0$ and $p \rightarrow \mu^+ \pi^0$ in 0.31 megaton-years exposure of the Super-Kamiokande Water Cherenkov Detector”, *Physical Review D* **95** (2017) 012004.
- [55] B. Fornal and B. Grinstein, “SU(5) Unification without Proton Decay”, *Physical Review Letters* **119** (2017) 241801.
- [56] N. Arkani-Hamed, A. G. Cohen, T. Gregoire, E. Katz, A. E. Nelson and J. G. Wacker, “The Minimal Moose for a Little Higgs”, *Journal of High Energy Physics* **2002** (2002) 021.
- [57] T. Han, H. E. Logan, B. McElrath and L.-T. Wang, “Phenomenology of the Little Higgs Model”, *Physical Review D* **67** (2003) 095004.
- [58] M. Cvetič and P. Langacker, “New Gauge Bosons from String Models”, *Modern Physics Letters A* **11** (1996) 1247.
- [59] A. Bilal, “Lectures on Anomalies”, arXiv note, 2008 [[arXiv:0802.0634](https://arxiv.org/abs/0802.0634)].
- [60] P. Langacker, “The Physics of Heavy Z’ Gauge Bosons”, *Reviews of Modern Physics* **81** (2009) 1199.

- [61] A. Rostomyan, “LNU, LNV and LFV at Belle II”, presentation at the 15th International Workshop on Tau Lepton Physics, Amsterdam, 2018.
- [62] W. Bertl, R. Engfer, E. Hermes, G. Kurz, T. Kozlowski, J. Kuth et al., “A search for μ -e conversion in muonic gold”, *The European Physical Journal C* **47** (2006) 337.
- [63] MEG Collaboration, “Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+ \gamma$ with the full dataset of the MEG experiment”, *The European Physical Journal C* **76** (2016) 434.
- [64] U. Bellgardt, G. Otter, R. Eichler, L. Felawka, C. Niebuhr, H. K. Walter et al., “Search for the decay $\mu^+ \rightarrow e^+ e^+ e^-$ ”, *Nuclear Physics B* **299** (1988) 1.
- [65] K. Hayasaka, “New Search for $\tau \rightarrow \mu \gamma$ and $\tau \rightarrow e \gamma$ Decays at Belle”, *Physics Letters B* **666** (2008) 16.
- [66] BABAR Collaboration, “Searches for Lepton Flavor Violation in the Decays $\tau \rightarrow e \gamma$ and $\tau \rightarrow \mu \gamma$ ”, *Physical Review Letters* **104** (2010) 021802.
- [67] K. Hayasaka, K. Inami and Y. Miyazaki, “Search for Lepton Flavor Violating Tau Decays into Three Leptons with 719 Million Produced $\tau^+ \tau^-$ Pairs”, *Physics Letters B* **687** (2010) 139.
- [68] BABAR Collaboration, “Limits on tau Lepton-Flavor Violating Decays in three charged leptons”, *Physical Review D* **81** (2010) 111101.
- [69] LHCb Collaboration, “Search for lepton-flavour-violating decays of Higgs-like bosons”, *The European Physical Journal C* **78** (2018) 1008.
- [70] CMS Collaboration, “Search for the lepton flavor violating decay $\tau \rightarrow 3\mu$ in proton-proton collisions at $\sqrt{s} = 13$ TeV”, submitted to *Journal of High Energy Physics*, 2020 [[arXiv:2007.05658](#)].
- [71] ATLAS Collaboration, “Probing lepton flavour violation via neutrinoless $\tau \rightarrow 3\mu$ decays with the ATLAS detector”, *The European Physical Journal C* **76** (2016) 232.
- [72] S. Bifani, S. Descotes-Genon, A. R. Vidal and M.-H. Schune, “Review of Lepton Universality tests in B decays”, *Journal of Physics G: Nuclear and Particle Physics* **46** (2019) 023001.
- [73] HFLAV Group, “Averages of b -hadron, c -hadron, and τ -lepton properties as of 2018 (2019 update)”, arXiv note, 2019 [[arXiv:1909.12524](#)].
- [74] A. de Gouvea and P. Vogel, “Lepton Flavor and Number Conservation, and Physics Beyond the Standard Model”, *Progress in Particle and Nuclear Physics* **71** (2013) 75.
- [75] ATLAS Collaboration, “Searches for lepton-flavour-violating decays of the Higgs boson in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”, *Physics Letters B* **800** (2020) 135069.
- [76] CMS Collaboration, “Search for lepton flavour violating decays of the Higgs boson to $\mu\tau$ and $e\tau$ in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of High Energy Physics* **2018** (2018) 1.
- [77] ATLAS Collaboration, “Search for lepton-flavor-violating decays of the Z boson into a τ lepton and a light lepton with the ATLAS detector”, *Physical Review D* **98** (2018) 092010.
- [78] CMS Collaboration, “Search for Lepton Flavor Violation in Z decays in pp collisions at $\sqrt{s} = 8$ TeV”, Physics Analysis Summary, CMS-PAS-EXO-13-005, 2015.
- [79] S. Nussinov, R. D. Peccei and X. M. Zhang, “Constraints on two-body lepton flavor violating decay processes”, *Physical Review D* **63** (2000) 016003.
- [80] CMS Collaboration, “Search for lepton flavour violating decays of a neutral heavy Higgs boson to $\mu\tau$ and $e\tau$ in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of High Energy Physics* **2020** (2020) 103.

- [81] ATLAS Collaboration, “Search for lepton-flavor violation in different-flavor, high-mass final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Physical Review D* **98** (2018) 092008.
- [82] CMS Collaboration, “Search for lepton-flavor violating decays of heavy resonances and quantum black holes to $e\mu$ final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of High Energy Physics* **2018** (2018) 73.
- [83] “The ISOLDE Radioactive Ion Beam facility”, website accessed in 2020-01-24 (<http://isolde.web.cern.ch>).
- [84] “The Antiproton Decelerator”, website accessed in 2020-01-24 (<https://home.cern/science/accelerators/antiproton-decelerator>).
- [85] “NA62”, website accessed in 2020-01-24 (<https://home.cern/science/experiments/na62>).
- [86] F. Marcastel, “CERN’s Accelerator Complex”, © 2013-2020 CERN (<https://cds.cern.ch/record/1621583>).
- [87] L. Evans and P. Bryant, “LHC Machine”, *Journal of Instrumentation* **3** (2008) .
- [88] ATLAS Collaboration, “The ATLAS Experiment at the CERN Large Hadron Collider”, *Journal of Instrumentation* **3** (2008) S08003.
- [89] ALICE Collaboration, “The ALICE experiment at the CERN LHC”, *Journal of Instrumentation* **3** (2008) S08002.
- [90] D. Boussard and T. P. R. Linnecar, “The LHC Superconducting RF System”, in *proceedings of the 1999 Cryogenic Engineering and International Cryogenic Materials Conference*, Montreal (Canada), 1999 (<https://cds.cern.ch/record/410377>).
- [91] “Accelerating: Radiofrequency cavities”, website accessed in 2020-01-29 (<https://home.cern/science/engineering/accelerating-radiofrequency-cavities>).
- [92] W. Höfle, T. P. R. Linnecar, R. Louwerse, J. Tückmantel, E. Gorbachev, I. N. Ivanov et al., “Transverse Damping Systems for the Future CERN LHC”, in *Proceedings of the 2001 Particle Accelerator Conference*, Chicago, pp. 1237–1239, 2001 (<https://accelconf.web.cern.ch/p01/PAPERS/TPAH004.PDF>).
- [93] CMS Collaboration, “The CMS experiment at the CERN LHC”, *Journal of Instrumentation* **3** (2008) S08004.
- [94] LHCb Collaboration, “The LHCb Detector at the LHC”, *Journal of Instrumentation* **3** (2008) S08005.
- [95] “2017 performance”, website accessed in 2020-01-29 (<https://lhc-commissioning.web.cern.ch/lhc-commissioning/performance/2017-performance.htm>).
- [96] R. Alemany, M. Lamont and S. Page, “LHC Modes”, LHC Project Document, LHC-OP-ES-0005, 2007.
- [97] J. Pedersen, G. Clark, K. D. Metzmacher, R. Losito, J. Gruber, H. O. Schönauer et al., “The PS complex as proton pre-injector for the LHC : design and implementation report”, CERN report, CERN-2000-003, 2000 (<https://cds.cern.ch/record/449242>).
- [98] A. Hilaire, V. Mertens and E. Weisse, “Beam transfer to and Injection into LHC”, in *Proceedings of the 6th European Particle Accelerator Conference*, Stockholm (Sweden), pp. 2117–2119, 1998 (<https://cds.cern.ch/record/364728>).
- [99] G. Arnison, A. Astbury, B. Aubert, C. Bacci, G. Bauer, A. Bézaguét et al., “Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s} = 540$ GeV”, *Physics Letters B* **122** (1983) 103.
- [100] M. Banner, R. Battiston, P. Bloch, F. Bonaudi, K. Borer, M. Borghini et al., “Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider”, *Physics Letters B* **122** (1983) 476.

- [101] G. Arnison, A. Astbury, B. Aubert, C. Bacci, G. Bauer, A. Bézaguet et al., “Experimental observation of lepton pairs of invariant mass around 95 GeV/ c^2 at the CERN SPS collider”, *Physics Letters B* **126** (1983) 398.
- [102] P. Bagnaia, M. Banner, R. Battiston, P. Bloch, F. Bonaudi, K. Borer et al., “Evidence for $Z_0 \rightarrow e^+e^-$ at the CERN pp collider”, *Physics Letters B* **129** (1983) 130.
- [103] CERN’s Engineering Data Management Service, “Interim summary report on the analysis of the 19 september 2008 incident at the LHC”, EDMS 973073, 2008 (https://edms.cern.ch/ui/file/973073/1/Report_on_080919_incident_at_LHC__2_.pdf).
- [104] LHC coordination, “LHC commissioning 2009”, website accessed in 2020-02-06, 2009 (<https://lhc-commissioning.web.cern.ch/lhc-commissioning/news-2009/LHC-2009.htm>).
- [105] LHC coordination, “2010 - Progress”, website accessed in 2020-02-06, 2010 (<https://lhc-commissioning.web.cern.ch/lhc-commissioning/progress2010.htm>).
- [106] P. Lujan, “CMS Luminosity - Public Results”, CMS Twiki accessed in 2020-02-07 (<https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>).
- [107] K. Borras, “First Run 2 results from CMS”, presentation at EPS-HEP Vienna, 2015.
- [108] CMS Collaboration, “The CMS trigger system”, *Journal of Instrumentation* **12** (2017) P01020.
- [109] G. Apollinari, O. Bruening, T. Nakamoto and L. Rossi, “High Luminosity Large Hadron Collider HL-LHC”, CERN Yellow Report, CERN-2015-005, 2017 [[arXiv:1705.08830](https://arxiv.org/abs/1705.08830)].
- [110] “LHC long term schedule”, website accessed in 2020-01-24 (<https://lhc-commissioning.web.cern.ch/lhc-commissioning/schedule/LHC-long-term.htm>).
- [111] “Detector | CMS Experiment”, website accessed in 2020-02-13 (<http://cms.cern/detector>).
- [112] K. Klein, “The Phase-1 upgrade of the CMS pixel detector”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **845** (2017) 101.
- [113] P. Lecoq, I. Dafinei, E. Auffray, M. Schneegans, M. V. Korzhik, O. V. Missevitch et al., “Lead tungstate (PbWO₄) scintillators for LHC EM calorimetry”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **365** (1995) 291.
- [114] K. Bell, R. Brown, D. Cockerill, P. Flower, P. Hobson, B. Kennedy et al., “Vacuum phototriodes for the CMS electromagnetic calorimeter endcap”, *IEEE Transactions on Nuclear Science* **51** (2004) 2284.
- [115] CMS Collaboration, “Energy calibration and resolution of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 7$ TeV”, *Journal of Instrumentation* **8** (2013) P09009.
- [116] CMS Collaboration, “CMS ECAL Performance in 2017”, CMS Detector Performance Summary, CMS-DP-2018/015, 2018 (<https://cds.cern.ch/record/2319285>).
- [117] CMS Collaboration, “CMS ECAL with Run2 data”, CMS Detector Performance Summary, CMS-DP-2019/005, 2019 (<https://cds.cern.ch/record/2668200/>).
- [118] D. Noonan, “Phase 1 upgrade of the CMS forward hadronic calorimeter”, *Journal of Instrumentation* **12** (2017) C02026.
- [119] A. Lobanov, “The CMS Outer HCAL SiPM Upgrade”, *Journal of Physics: Conference Series* **587** (2015) 012005.

- [120] “HCAL Completes Phase 1 Upgrades on the CMS Detector”, website accessed in 2020-02-26 (<https://cms.cern/news/hcal-completes-phase-1-upgrades-cms-detector>).
- [121] CMS Collaboration, “Calibration of the CMS hadron calorimeters using proton-proton collision data at $\sqrt{s} = 13$ TeV”, *Journal of Instrumentation* **15** (2020) P05002 [[arXiv:1910.00079](https://arxiv.org/abs/1910.00079)].
- [122] V. D. Elvira, “Measurement of the Pion Energy Response and Resolution in the CMS HCAL Test Beam 2002 Experiment”, CMS Note, CMS-NOTE-2004-020, 2004 (<https://cds.cern.ch/record/800406>).
- [123] J. G. Layter, “The CMS muon project : Technical Design Report”, Technical Design Report CMS, CMS-TDR-3, 1997 (<https://cds.cern.ch/record/343814>).
- [124] M. S. Kim, Y. Ban, J. Cai, Q. Li, S. Liu, S. Qian et al., “Web-based monitoring tools for Resistive Plate Chambers in the CMS experiment at CERN”, *Journal of Instrumentation* **9** (2014) .
- [125] M. Abbrescia, A. Colaleo, G. Iaselli, F. Loddo, M. Maggi, B. Marangelli et al., “The simulation of resistive plate chambers in avalanche mode: charge spectra and efficiency”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **431** (1999) 413.
- [126] R. Cardarelli, A. Di Ciaccio and R. Santonico, “Performance of a resistive plate chamber operating with pure CF₃Br”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **333** (1993) 399.
- [127] CMS Collaboration, “Performance of the reconstruction and identification of high-momentum muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of Instrumentation* **15** (2020) P02027.
- [128] D. Abbaneo, M. Abbas, M. Abbrescia, M. A. Akl, O. Aboamer, D. Acosta et al., “Charged particle detection performance of Gas Electron Multiplier (GEM) detectors for the upgrade of CMS endcap muon system at the CERN LHC”, in *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, San Diego (USA), pp. 1–4, 2015.
- [129] I. Vai, “Commissioning and performance of the GE1/1 slice test detectors”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **936** (2019) 439.
- [130] L. Thomas, “CMS High Level Trigger performance at 13 TeV”, in *Proceedings of The 39th International Conference on High Energy Physics*, vol. 340, Seoul, p. 226, 2019 (<https://pos.sissa.it/340/226>).
- [131] A. Tapper and D. Acosta, “CMS Technical Design Report for the Level-1 Trigger Upgrade”, Technical Design Report CMS, CMS-TDR-12, 2013 (<https://cds.cern.ch/record/1556311>).
- [132] CMS Collaboration, “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”, submitted to *Journal of Instrumentation*, 2020 [[arXiv:2006.10165](https://arxiv.org/abs/2006.10165)].
- [133] W. H. Smith, “Calorimeter Trigger Algorithms”, website accessed in 2020-03-05 (<https://pages.hep.wisc.edu/~wsmith/cms/calgeo/>).
- [134] F. Thiant, Y. Geerebaert, F. Magniette, T. Romanteau, A. Zabi, A. Zghiche et al., “New development in the CMS ECAL Level-1 trigger system to meet the challenges of LHC Run 2”, in *Proceedings of Topical Workshop on Electronics for Particle Physics*, vol. 343, Antwerp (Belgium), p. 052, 2019 (<https://pos.sissa.it/343/052/>).

- [135] P. Arce, M. Bellato, M. Benettoni, A. Benvenuti, D. Bonacorsi, M. Bontenackels et al., “Bunched beam test of the CMS drift tubes local muon trigger”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **534** (2004) 441.
- [136] A. Triossi, M. Bellato, J. M. C. Ruiz, J. Ero, C. F. Bedoya, G. Flouris et al., “The CMS Barrel Muon trigger upgrade”, *Journal of Instrumentation* **12** (2017) C01095.
- [137] D. Acosta, A. Brinkerhoff, E. Busch, A. Carnes, I. Furic, S. Gleyzer et al., “Boosted Decision Trees in the Level-1 Muon Endcap Trigger at CMS”, *Journal of Physics: Conference Series* **1085** (2018) 042042.
- [138] N. Daci, “CMS Trigger Performances”, presentation at the 7th International Conference on New Frontiers in Physics, Orthodox Academy of Crete (Greece), 2018 (<https://indico.cern.ch/event/663474/contributions/3061225>).
- [139] S. Mukherjee, “Data Scouting : A New Trigger Paradigm”, in *Proceedings of the Fifth Annual LHCP*, Shanghai (China), 2017, [[arXiv:1708.06925](https://arxiv.org/abs/1708.06925)].
- [140] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *Journal of Instrumentation* **12** (2017) P10003.
- [141] D. Barney, “Cms Slice”, Education and Outreach document, CMS-OUTREACH-2018-017, 2018 (<https://cds.cern.ch/record/2628641>).
- [142] CMS Collaboration, “Description and performance of track and primary-vertex reconstruction with the CMS tracker”, *Journal of Instrumentation* **9** (2014) P10009.
- [143] R. Frühwirth, “Application of Kalman filtering to track and vertex fitting”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **262** (1987) 444.
- [144] CMS Collaboration, “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of Instrumentation* (2018) 06015.
- [145] G. Abbiendi, I. Bachiller, C. Battilana, S.S. Chhibra, B. De La Cruz, F. Errico et al., “Performance of high-pt muons collected with CMS at $\sqrt{s} = 13$ TeV during proton-proton collisions”, CMS Draft Analysis Note, AN-2018/008, 2019.
- [146] W. Adam, R. Frühwirth, A. Strandlie and T. Todorov, “Reconstruction of electrons with the Gaussian-sum filter in the CMS tracker at LHC”, *Journal of Physics G: Nuclear and Particle Physics* **31** (2005) N9.
- [147] CMS Collaboration, “Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *Journal of Instrumentation* **10** (2015) P06005.
- [148] F. Wenxing, *Search for new physics in dilepton final states at the CMS experiment*, Ph.D. thesis, Université libre de Bruxelles, 2019.
- [149] M. Cacciari, G. P. Salam and G. Soyez, “The anti-kt jet clustering algorithm”, *Journal of High Energy Physics* **2008** (2008) 063.
- [150] CMS Collaboration, “Pileup mitigation at CMS in 13 TeV data”, *Journal of Instrumentation* **15** (2020) P09018.
- [151] M. Cacciari and G. P. Salam, “Pileup subtraction using jet areas”, *Physics Letters B* **659** (2008) 119.
- [152] CMS Collaboration, “Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS”, *Journal of Instrumentation* **6** (2011) P11002.
- [153] CMS Collaboration, “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV”, *Journal of Instrumentation* **13** (2018) P05011.

- [154] CMS Collaboration, “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector”, *Journal of Instrumentation* **14** (2019) P07004.
- [155] M. Tanabashi et al. (Particle Data Group), “The Review of Particle Physics (2019)”, *Physical Review D* **98** (2018 and 2019 update) 030001.
- [156] CMS Collaboration, “Performance of tau-lepton reconstruction and identification in CMS”, *Journal of Instrumentation* **7** (2012) P01001.
- [157] CMS Collaboration, “Performance of reconstruction and identification of tau leptons in their decays to hadrons and tau neutrino in LHC Run-2”, CMS Physics Analysis Summary, CMS-PAS-TAU-16-002, 2016 (<https://cds.cern.ch/record/2196972>).
- [158] CMS Collaboration, “Reconstruction and identification of tau lepton decays to hadrons and tau neutrino at CMS”, *Journal of Instrumentation* **11** (2016) P01019.
- [159] C. Böser, S. Fink and S. Röcker, “Introduction to Boosted Decision Trees”, presentation at the KSETA Doktoranden Workshop, Waldhotel Zollernblick (Germany), 2014.
- [160] Glosser.ca, “Artificial neural network with layer coloring”, file retrieved from Wikimedia (https://commons.wikimedia.org/wiki/File:Colored_neural_network.svg).
- [161] K. Androsov, “Identification of tau leptons using Deep Learning techniques at CMS”, in *Proceedings of the 27th International Symposium on Nuclear Electronics and Computing*, Becici, Budva (Montenegro), 2019 (<https://cds.cern.ch/record/2713735>).
- [162] CMS Collaboration, “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV”, *Journal of Instrumentation* **13** (2018) P10005.
- [163] “Facts and figures about the LHC”, website accessed in 2020-05-11 (<https://home.cern/resources/faqs/facts-and-figures-about-lhc>).
- [164] A. Buckley, J. Butterworth, S. Gieseke, D. Grellscheid, S. Hoche, H. Hoeth et al., “General-purpose event generators for LHC physics”, *Physics Reports* **504** (2011) 145.
- [165] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce Dubois, M. Asai et al., “Geant4 developments and applications”, *IEEE Transactions on Nuclear Science* **53** (2006) 270.
- [166] J. Allison, K. Amako, J. Apostolakis, P. Arce, M. Asai, T. Aso et al., “Recent developments in Geant4”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **835** (2016) 186.
- [167] NNPDF Collaboration, “Parton distributions for the LHC Run II”, *Journal of High Energy Physics* **2015** (2015) 40.
- [168] T. Sjöstrand, S. Mrenna and P. Skands, “PYTHIA 6.4 physics and manual”, *Journal of High Energy Physics* **2006** (2006) 026.
- [169] T. Sjöstrand, S. Mrenna and P. Skands, “A Brief Introduction to PYTHIA 8.1”, *Computer Physics Communications* **178** (2008) 852.
- [170] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten et al., “An Introduction to PYTHIA 8.2”, *Computer Physics Communications* **191** (2015) 159.
- [171] CMS Collaboration, “Event generator tunes obtained from underlying event and multiparton scattering measurements”, *The European Physical Journal C* **76** (2016) 155.
- [172] A. Belyaev, N. D. Christensen and A. Pukhov, “CalcHEP 3.4 for collider physics within and beyond the Standard Model”, *Computer Physics Communications* **184** (2013) 1729.

- [173] D. M. Gingrich, “Monte Carlo event generator for black hole production and decay in proton-proton collisions”, *Computer Physics Communications* **181** (2010) 1917.
- [174] NNPDF Collaboration, “Parton distributions from high-precision collider data”, *The European Physical Journal C* **77** (2017) 663 [[arXiv:1706.00428](https://arxiv.org/abs/1706.00428)].
- [175] J. Pumplin, D. R. Stump, J. Huston, H. L. Lai, P. Nadolsky and W. K. Tung, “New Generation of Parton Distributions with Uncertainties from Global QCD Analysis”, *Journal of High Energy Physics* **2002** (2002) 012.
- [176] S. Alioli, P. Nason, C. Oleari and E. Re, “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”, *Journal of High Energy Physics* **2010** (2010) 43.
- [177] S. Kallweit, J. M. Lindert, S. Pozzorini and M. Schonherr, “NLO QCD+EW predictions for $2\ell 2\nu$ diboson signatures at the LHC”, *Journal of High Energy Physics* **2017** (2017) 120.
- [178] M. Czakon, D. Heymes, A. Mitov, D. Pagani, I. Tsinikos and M. Zaro, “Top-pair production at the LHC through NNLO QCD and NLO EW”, *Journal of High Energy Physics* **2017** (2017) 186.
- [179] R. Frederix and S. Frixione, “Merging meets matching in MC@NLO”, *Journal of High Energy Physics* **2012** (2012) 61.
- [180] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, *Journal of High Energy Physics* **2014** (2014) 79.
- [181] F. Cascioli, T. Gehrmann, M. Grazzini, S. Kallweit, P. Maierhöfer, A. von Manteuffel et al., “ZZ production at hadron colliders in NNLO QCD”, *Physics Letters B* **735** (2014) 311.
- [182] F. Campanario, C. Englert, S. Kallweit, M. Spannowsky and D. Zeppenfeld, “NLO QCD corrections to WZ+jet production with leptonic decays”, *Journal of High Energy Physics* **2010** (2010) 76.
- [183] A. Andonov, A. Arbuzov, S. B. P. Christova, V. Kolesnikov, G. Nanava and R. Sadykov, “NLO QCD corrections to Drell-Yan processes in the SANC framework”, *Phys. Atom. Nuclei* **73** (2010) 1761.
- [184] E. L. Berger, J. Gao, C.-P. Yuan and H. X. Zhu, “NNLO QCD Corrections to t-channel Single Top-Quark Production and Decay”, *Physical Review D* **94** (2016) 071501.
- [185] CMS Collaboration, “Muon HLT Performance with 2018 Data”, CMS Performance Note, CMS-DP-2018/034, 2018 (<https://cds.cern.ch/record/2627469>).
- [186] CMS Collaboration, “Search for a narrow resonance in high-mass dilepton final states in proton-proton collisions using 140 fb⁻¹ of data at $\sqrt{s}=13$ TeV”, CMS Physics Analysis Summary, CMS-PAS-EXO-19-019, 2019 (<https://cds.cern.ch/record/2684757>).
- [187] CMS Collaboration, “Search for narrow resonances in dilepton mass spectra in proton-proton collisions at $\sqrt{s} = 13$ TeV and combination with 8 TeV data”, *Physics Letters B* **768** (2017) 57.
- [188] CMS Collaboration, “Search for high mass resonances in dielectron final state”, CMS Physics Analysis Summary, CMS-PAS-EXO-18-006, 2018 (<https://cds.cern.ch/record/2308270>).
- [189] B. Cousins, “Probability Density Functions for Positive Nuisance Parameters”, CMS Statistics Committee report, 2010 (http://www.physics.ucla.edu/~cousins/stats/cousins_lognormal_prior.pdf).

- [190] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht et al., “LHAPDF6: parton density access in the LHC precision era”, *The European Physical Journal C* **75** (2015) 132.
- [191] G. Abbiendi, A. Alfonsi, D. Bourilkov and S.S. Chhibra, “Search for High-Mass Resonances Decaying to Muon Pairs in pp Collisions at $\sqrt{s}=13$ TeV with the full 2016 dataset of 37 fb^{-1} and combination with 2015 result”, CMS Draft Analysis Note, CMS-AN-2016/391, 2016.
- [192] B. Clerbaux, D. Cockerill, G. Daskalakis, M. Dragnet, Sh. Elgammal, W. Fang et al., “Search for high-mass resonances in the di-electron finalstate with 2018 data”, Analysis Note CMS AN-18-253, 2018.
- [193] R. D. Cousins, “Why isn’t every physicist a Bayesian?”, *American Journal of Physics* **63** (1995) 398.
- [194] R. D. Cousins, “Lectures on Statistics in Theory: Prelude to Statistics in Practice”, lecture summary, 2018 [[arXiv:1807.05996](https://arxiv.org/abs/1807.05996)].
- [195] CMS Collaboration, “Combine”, CMS Analysis Github page (<http://cms-analysis.github.io/HiggsAnalysis-CombinedLimit/>).
- [196] CMS Collaboration, “Search for heavy resonances decaying to tau lepton pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of High Energy Physics* **2017** (2017) 48.
- [197] ATLAS Collaboration, “Search for additional heavy neutral Higgs and gauge bosons in the ditau final state produced in 36 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Journal of High Energy Physics* **2018** (2018) 55.
- [198] ATLAS Collaboration, “Search for high-mass dilepton resonances using 139 fb^{-1} of pp collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Physics Letters B* **796** (2019) 68.
- [199] M. Lindner, M. Platscher and F. S. Queiroz, “A Call for New Physics : The Muon Anomalous Magnetic Moment and Lepton Flavor Violation”, arXiv note, 2017 [[arXiv:1610.06587](https://arxiv.org/abs/1610.06587)].
- [200] S. Antusch, A. Hammad and A. Rashed, “Probing Z mediated charged lepton flavor violation with taus at the LHeC”, *Physics Letters B* **810** (2020) 135796.