Search for Charged Higgs Bosons Decaying to a Tau Lepton and a Neutrino with the CMS Experiment

Santeri Laurila

Search for
Charged Higgs Bosons
Decaying to a
Tau Lepton and a Neutrino
with the CMS Experiment

Santeri Laurila

Helsinki Institute of Physics
University of Helsinki

ACADEMIC DISSERTATION

To be presented, with the permission of the Faculty of Science of the University of Helsinki,
for public criticism in the auditorium CK112 at Exactum, Gustaf Hällströmin katu 2 B,
Helsinki, on the 18th of October 2019 at 12 o'clock

Helsinki 2019

Acknowledgments

Science is fundamentally a collective endeavor. This work is built upon the foundations created by generations of scientists before me, and made possible by all the hard work from thousands of past and present members of the CMS Collaboration. I want to thank all the taxpayers of CERN member states for maintaining such an exceptional research environment, which enables strong international cooperation for a shared, ambitious goal—a deeper understanding of nature. I believe that in addition to their scientific value, large international research facilities like CERN also serve an important societal function, as they demonstrate the power of peaceful collaboration based on the values of openness and diversity.

There are a number of people that I wish to thank personally. First and foremost I thank my thesis supervisor Prof. Paula Eerola and my thesis advisor Doc. Sami Lehti. They have always encouraged me and given me the space to grow into an independent scientist, while providing support and constructive feedback when I have needed it. Even as a vice rector of the university, Paula has found the time to supervise my work, sharing her broad knowledge and long experience in the field. Thanks to Sami's great expertise in Higgs physics and data analysis, his advice has been very helpful and instructive. I also thank Dr. Lauri Wendland for all his help and advice in the early phase of this thesis project.

While preparing the results presented in this thesis, I worked closely with Dr. Jan Eysermans who was in charge of a parallel analysis, targeting the leptonic final states of the same search channel. I wish to thank Jan for the efficient and enjoyable cooperation. I would also like to thank Dr. Andrea Marini and his colleagues at the Massachusetts Institute of Technology, as well as Dr. Alexandros Attikis and his colleagues at the University of Cyprus, for their contributions in the data analysis and for useful discussions. I express my gratitude to all my CMS colleagues who participated in reviewing the analysis and gave feedback on the article manuscripts.

Abstract

The Standard Model of particle physics is the most successful and precise theoretical description of fundamental physics so far. The discovery of the Higgs boson in 2012 by the ATLAS and CMS experiments at the LHC provided strong evidence for the Englert-Brout-Higgs-Guralnik-Hagen-Kibble mechanism, explaining how elementary particles gain their masses in the Standard Model. However, the Standard Model is known to be an incomplete description of nature, as it cannot explain the origin of dark matter, neutrino masses or the observed matter–antimatter asymmetry. Therefore more general models with an extended Higgs sector are actively being studied. Models with at least two Higgs doublets predict the existence of electrically charged Higgs bosons. The observation of charged Higgs bosons would provide direct evidence for new physics and guide the way towards a more comprehensive theory.

In this thesis, a search is presented for charged Higgs bosons decaying into a tau lepton and a neutrino, based on proton-proton collision events recorded by the CMS experiment in 2016 at a center-of-mass energy of 13 TeV. The amount of data corresponds to an integrated luminosity of $35.9\,\mathrm{fb}^{-1}$. The search targets the hadronic final state with a hadronically decaying tau lepton, missing momentum due to neutrinos, and additional jets from an associated top quark decay.

This analysis contains multiple methodological improvements with respect to the previous CMS results on the same search channel. The particle identification algorithms and selection criteria are optimized for good performance under challenging luminosity conditions. Categorization of events based on tau lepton helicity is used to enhance sensitivity. The background from events with jets misidentified as tau leptons is estimated from data, whereas the background from genuine-tau events is estimated from simulation. This thesis also presents a new version of the tau embedding method, which allows the estimation of the genuine-tau background using single-muon events.

The transverse mass of the tau-neutrino system is reconstructed. As the data agree with the background-only hypothesis, upper limits are derived for the charged Higgs boson production rate. The search covers signal hypotheses from 80 GeV to 3 TeV, and for the first time in CMS, the hypotheses with the charged Higgs boson mass close to the top quark mass are scanned. For maximal signal sensitivity, the results are combined with those from the leptonic final states of the same search channel. The combined result is interpreted in the context of the Minimal Supersymmetric Standard Model.

Tiivistelmä

Hiukkasfysiikan standardimalli on toistaiseksi tarkin teoria perustavista luonnonlaeista. Higgsin bosonin havaitseminen vuonna 2012 LHC-törmäyttimen ATLAS- ja CMS-koeasemilla tarjosi vahvaa näyttöä niin sanotun Higgsin mekanismin puolesta. Tämä mekanismi selittää, kuinka alkeishiukkasten massat syntyvät standardimallissa. Standardimalli ei kuitenkaan ole täydellinen kuvaus luonnosta, sillä se ei selitä pimeän aineen syntyperää, neutriinojen massoja tai havaittua aineen ja antiaineen epäsymmetriaa. Siksi yleisempiä malleja, joissa on useampia Higgsin kenttiä, tutkitaan tiiviisti. Mallit, joissa on vähintään kaksi Higgs-duplettikenttää, ennustavat sähkövarauksellisten Higgsin bosonien olemassaolon. Sähkövarauksellisten Higgsin bosonien havaitseminen tarjoaisi suoraa todistusaineistoa standardimallin ulkopuolisesta fysiikasta ja viitoittaisi tietä kohti laajempaa teoriaa.

Tässä väitöskirjassa etsitään sähkövarauksellisia Higgsin bosoneita, jotka hajoavat tau-leptoniksi ja neutriinoksi. Etsinnässä analysoidaan CMS-koeaseman vuonna 2016 mittaamia 13 TeV:n massakeskipiste-energian protoni–protoni-törmäyksiä. Datan määrä vastaa 35.9 fb^{-1}:n integroitua luminositeettia. Tähtäimessä on hadroninen lopputila, jossa tau-leptoni hajoaa hadronisesti, top-kvarkin hajoamisesta syntyy hadronisia hiukkasryöppyjä (jettejä) ja jossa neutriinot havaitaan epäsuorasti mitattujen kohtisuorien liikemäärien epätasapainona.

Työssä käytetyt analyysimenetelmät ovat kehittyneempiä verrattuna aiempiin CMS-kollaboraation tuloksiin samasta etsintäkanavasta. Hiukkasten tunnistamiseen ja törmäystapahtumien valitsemiseen käytetyt algoritmit on optimoitu korkean luminositeetin olosuhteisiin. Tilastollisessa analyysissä törmäystapahtumat kategorisoidaan tau-leptonin helisiteettiin perustuvalla menetelmällä. Taustatapahtumat, joissa jetti on virheellisesti tunnistettu tau-leptoniksi, arvioidaan datasta. Oikeita tau-leptoneita sisältävät taustatapahtumat puolestaan arvioidaan simulaatiolla. Tässä väitöskirjassa esitellään myös uusi versio menetelmästä, jolla tämä tausta voidaan arvioida myoneja sisältävistä törmäystapahtumista.

Tilastollinen analyysi perustuu tau–neutriino-systeemin kohtisuoraan massaan, jonka avulla johdetaan yläraja sähkövarauksellisten Higgsin bosonien tuottotodennäköisyydelle. Etsintä kattaa massa-alueen 80 GeV:n ja 3 TeV:n välillä. Tämä on ensimmäinen CMS-analyysi, jossa etsintään sisältyy myös massa-alue top-kvarkin massan läheisyydessä. Tulokset yhdistetään saman etsintäkanavan leptonisen lopputilan tuloksiin ja tulkitaan käyttäen minimaalista supersymmetristä standardimallia.

"We were always
After something
We were always
Chasing something
Until moving
Becomes everything
Until moving
Just becomes everything"
— New Model Army: After Something

Publications

I. CMS Collaboration (incl. S. Laurila), "Search for charged Higgs bosons with the $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ decay channel in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **07** (2019) 142, 2019. doi:10.1007/JHEP07(2019)142 (`arxiv.org/abs/1903.04560`)

II. CMS Collaboration (incl. S. Laurila), "Search for charged Higgs bosons with the $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ decay channel in the fully hadronic final state at $\sqrt{s} = 13$ TeV", *CMS Physics Analysis Summary* CMS-PAS-16-031, 2016. (`cds.cern.ch/record/2223865`)

III. S. Laurila for CMS Collaboration, "Searches for charged Higgs bosons at CMS", Proceedings of Science ICHEP2018 155, 2018.

IV. S. Laurila for CMS Collaboration, "Recent results on 2HDM charged Higgs boson searches in CMS", Proceedings of Science CHARGED2016 008, 2017.

Author's contribution

The main results of this thesis are published in publication I. In this paper, the author had a leading role in the analysis work, carrying out the analysis in the fully hadronic final state, acting as the contact person of the analysis inside CMS Collaboration, as well as writing and editing the paper submitted to *Journal of High Energy Physics*. The analysis is detailed in Part III of this thesis.

A preliminary version of the analysis with a partial data set was presented in publication II. For this result, the author was responsible for the optimization of event selection, implementation of the systematic uncertainties and the final statistical analysis of the data.

The author presented the charged Higgs boson search results by the CMS Collaboration at the *XXXIX International Conference on High Energy Physics* (*ICHEP2018*) and at the *Sixth International Workshop on Prospects for Charged Higgs Discovery at Colliders* (*CHARGED16*), and wrote the corresponding conference proceedings III and IV.

The author participated in the data taking effort in the CMS experiment with several contributions to monitoring and data certification for the upgraded Level-1 trigger system, including the *Level-1 trigger offline data certification co-coordinator* responsibility in 2018. Therefore Section 6 of this thesis is dedicated to the upgraded CMS trigger system and its operation.

Contents

Acronyms and symbols

BDT	Boosted decicion tree
BSM	Beyond the Standard Model
CERN	European Organization for Nuclear Research
C.L.	Confidence level
CKM	Cabibbo–Kobayashi–Maskawa [matrix]
CMS	Compact Muon Solenoid
CSC	Cathode strip chamber [muon station]
CTF	Combinatorial track finder [algorithm]
DAQ	Data acquisition
DT	Drift tube [muon station]
DQM	Data quality monitoring
ECAL	Electromagnetic calorimeter
EWK	Electroweak
FCNC	Flavor-changing neutral current
FPGA	Field programmable gate array
GEM	Gas electron multiplier [muon station]
GSF	Gaussian-sum filter
HB	Hadron barrel calorimeter
HCAL	Hadron calorimeter
HE	Hadron endcap calorimeter
HLT	High-level trigger
HF	Hadron forward calorimeter
HPS	Hadron-plus-strips [algorithm]
IP	Impact parameter
L1A	Level-1 Accept [signal]
L1T	Level-1 trigger
LEP	Large Electron Positron Collider
LHC	Large Hadron Collider
LO	Leading order
MSSM	Minimal Supersymmetric Standard Model
NLO	Next-to-leading order
PDF	Parton density function
pdf	Probability density function
PF	Particle flow [algorithm]
PV	Primary vertex
RPC	Resistive plate chamber
SM	Standard Model
SV	Secondary vertex
TEC	Tracker endcap
TIB	Tracker inner barrel
TOB	Tracker outer barrel
TT	Trigger tower [in calorimeters]

QCD	Quantum chromodynamics
2HDM	Two-Higgs-doublet model
$\mathcal{B}$	Branching fraction
c	Speed of light [in vacuum]
η	Pseudorapidity
ϕ	Azimuthal angle
Γ	Decay width
h.c.	Hermitian conjugate
$\mathcal{L}$	Lagrangian density *or* instantaneous luminosity *or* likelihood function
ℓ	Charged lepton
m	Invariant mass
m_T	Transverse mass
$\vec{p}_\mathrm{T}$	Transverse momentum
$\vec{p}_\mathrm{T}^{\,\mathrm{miss}}$	Missing transverse momentum
ΔR	Distance in η–ϕ plane
σ	Cross section
τ_h	Hadronically decaying tau lepton

Conventions

Charge conjugation is implied throughout the thesis, e.g. by "electrons" we mean both electrons and positrons. For clarity, generic notation is used when possible: for example, $\mathrm{H}^+ \to \tau^+ \nu_\tau$ and $\mathrm{H}^- \to \tau^- \bar{\nu}_\tau$ processes are commonly denoted as $\mathrm{H}^\pm \to \tau \nu_\tau$. Similarly, $\mathrm{H}^+ \to \mathrm{t\bar{b}}$ and $\mathrm{H}^- \to \mathrm{\bar{t}b}$ are denoted as $\mathrm{H}^\pm \to \mathrm{tb}$.

The CMS coordinate system is explained in Section 5. In this coordinate system, the four-momentum $(E, \vec{p})$ of an object is convenient to parametrize as $(p_\mathrm{T}, \eta, \phi, m)$. In this thesis, these variables are used to describe particle kinematics.

Natural units with $\hbar = c = 1$ are used throughout the thesis. In natural units, energies E, momenta p and masses m are related via $E^2 = p^2 + m^2$, so they can all be expressed in units of electron volts (eV). Typically giga electron volts (GeV, 10^9 eV) or tera electron volts (TeV, 10^{12} eV) are used.

Transvese energy is defined as $E_\mathrm{T}^2 = p_\mathrm{T}^2 + m^2$. For objects with $p_\mathrm{T} \gg m$, $E_\mathrm{T} \approx p_\mathrm{T}$, in which case we do not distinguish between the notions of transverse momentum and transverse energy.

The symbol $\vec{p}_\mathrm{T}^{\,\mathrm{miss}}$ refers to the type-I corrected missing transverse momentum. For historical reasons, its magnitude ($p_\mathrm{T}^\mathrm{miss}$) is often referred to as missing transverse energy, but here also the magnitude is called missing transverse momentum.

Chapter 1

Introduction

Particle physics is a branch of science that studies Nature at its most fundamental level, understood in the current theory as elementary particles and their interactions. The word *physics* has its origins in Ancient Greek where *physis* means nature, whereas the Latin word *particula* refers to small parts. In quantum physics, small distances correspond to large energy scales, so studying the smallest constituents of nature requires theoretical understanding of high-energy processes and advanced technology to produce them in a laboratory. Thus the field of particle physics is also often called *high-energy physics*.

The theoretical and experimental efforts by generations of scientists have allowed us to zoom into the structure of matter with higher and higher precision, and to find new fascinating structures at every level. Now we understand that chemistry is driven by the properties of atoms, and that atoms contain electrons and atomic nuclei made of protons and neutrons, which in turn are made of quarks and gluons.

In the 1960s and early 1970s, the current understanding of fundamental physics was formulated as a unified quantum field theory known as the *Standard Model* (SM) of particle physics [1–6]. It is one of the most rigorous and precise scientific theories ever created, and it has passed innumerable experimental tests over the past decades. The SM predicted the existence of several elementary particles, such as $W^{\pm}$ and Z bosons, gluons and the top quark, before they were experimentally discovered.

The last missing piece of the SM pending experimental confirmation was the existence of a *Higgs boson*. In the SM, elementary particles gain their masses by interacting with a field known as the Higgs field, manifesting itself as Higgs bosons. The Englert-Brout-Higgs-Guralnik-Hagen-Kibble mechanism [7–12] (or Higgs mechanism, for short)

that makes this possible via a spontaneous breaking of electroweak symmetry was predicted by three independent groups in 1964.

Massive elementary particles such as Higgs bosons can be produced in energetic particle collisions, converting the energy of the colliding particles (E), which is mostly kinetic energy, into mass (m) according to Einstein's famous equation $E = mc^2$. This is the principle behind the Large Hadron Collider (LHC), the largest machine ever built [13]. In 2012, the ATLAS and CMS experiments at the LHC discovered a new particle with a mass of approximately 125 GeV [14, 15], which was later confirmed to be a Higgs boson [16]. The discovery completed the era of experimental searches for new particles guided by the SM.

While the SM is one of the most successful theories developed this far, it suffers from both experimental and theoretical shortcomings. Current cosmological observations suggest that the particles described by the SM constitute only some 16% of the total mass of the matter in the universe. The rest is *dark matter* which might have escaped optical observation, because it does not interact electromagnetically. In addition to lacking a particle suitable as a dark matter candidate, the SM fails to explain the observed imbalance between matter and antimatter in the universe around us, a prerequisite for our very existence. Furthermore, the discovery of neutrino oscillations in 1998 proved that these weakly interacting particles are not exactly massless, as assumed in the SM. On the theoretical side, the SM lacks explanations for some of its peculiar features, such as the hierarchy problem related to the radiative corrections to the Higgs boson mass, and the fact that it contains three generations of particles with similar properties but different masses.

All of this suggests that the SM is not a complete description of nature, but rather a low-energy approximation of a more general theory. Many candidates for this wider *beyond the Standard Model* (BSM) theory have been proposed. Almost all of these theories predict an extended Higgs sector, with a spectrum of Higgs bosons with different masses, charges, and other properties. These models are constrained, but not excluded, by the measured properties of the 125 GeV boson. Two-Higgs-doublet models (2HDMs) predict five different Higgs bosons: two neutral CP-even particles h and H (with $m_\mathrm{h} \leq m_\mathrm{H}$), one neutral CP-odd particle A, and two charged Higgs bosons $H^\pm$ [17, 18]. The observation of additional Higgs bosons would provide direct evidence for the existence of BSM physics, and guide the theoretical developments towards a new theory.

In this thesis, a search for charged Higgs bosons is performed, based on proton-proton collisions provided by the LHC at a center-of-mass energy of 13 TeV and collected by the CMS experiment in 2016. The amount of data corresponds to an integrated luminosity of $35.9\,\text{fb}^{-1}$. The search targets collision events where a charged Higgs boson ($H^{\pm}$) is produced and immediately decays into a tau lepton (τ) and a neutrino (ν). In many models, this is the experimentally most sensitive channel for the observation of the $H^{\pm}$.

Thesis in a nutshell

The Standard Model is reviewed in Chapter 2. Its shortcomings are discussed, and the extended models that motivate the searches for charged Higgs bosons are presented in Chapter 3.

Direct searches for $H^{\pm}$ have been performed at the CERN LEP [19], at the Fermilab Tevatron [20,21], and by the LHC experiments. As several BSM models contain charged Higgs bosons, and probabilities for their different production and decay modes depend on the details of the theory, the LHC experiments have broad $H^{\pm}$ search programs. The ATLAS and CMS Collaborations have covered several $H^{\pm}$ decay channels, such as $\tau^{\pm}\nu_{\tau}$, tb, cs, cb and $W^{\pm}$ Z in their previous searches at center-of-mass energies of 7, 8, or 13 TeV [22–33]. Additionally, the results from ATLAS and CMS searches for additional neutral Higgs bosons, electroweak precision measurements, and flavor physics, can be interpreted as constraints in the 2HDM parameter space. These results and their implications for $H^{\pm}$ searches are reviewed in Chapter 3.

The LHC, the most powerful particle collider ever built, is situated in a 27-kilometer circular tunnel near Geneva, Switzerland. It started operating in 2008 and is foreseen to continue running for decades, with frequent upgrades to further improve its performance. The LHC and basic concepts of collider physics are introduced in Chapter 4.

Four experiments, each the size of a building, are situated underground along the accelerator tunnel to observe and register the collisions produced by the LHC. One of them is the Compact Muon Solenoid (CMS), a general-purpose experiment composed of cylindrical layers of tracking detectors, calorimeters and muon detectors. Protons accelerated by the LHC are set to collide at the center of CMS, creating new particles that produce electric signals as they traverse the detector. The design and working principles of this complicated apparatus are introduced in Chapter 5. The CMS

experiment is operated by and its data is analyzed by a collaboration of some 1000 engineers and 3000 scientists working in 231 member institutes around the world. Finnish institutes and physicists have been actively involved in the experiment since the beginning of the project.

Inside the CMS detector, proton bunches cross 40 million times per second, with tens of proton-proton interactions at each crossing *event*. Thus approximately one billion proton-proton interactions take place every second. As CMS produces approximately one megabyte of data per event, the readout, storage and analysis of all events would be impossible. On the other hand, only a small fraction of the events are of special interest. For example, the average rate of 125 GeV Higgs boson production is only once per second.

Thus a sophisticated *trigger system* is needed to select approximately a few hundred most interesting collision events per second, to be stored for further analysis. The initial decision to keep or reject an event needs to be done within a few microseconds and it is irreversible. As the trigger system is a critical component in the experiment, it needs to be constantly monitored. The CMS trigger system and the procedures to ensure its efficient and reliable functioning are described in Chapter 6.

Even after the reduction of data flow with the trigger system, storage and post-processing of the data is a massive computational challenge. The CERN data center hosts hundreds of petabytes of collision data, distributed around the world via the LHC worldwide computing grid. The recorded data are accompanied by large amounts of simulated events, produced using Monte Carlo simulation techniques to model the physics processes and the detector response. The computational methods for data processing and simulation methods are discussed in Chapter 7.

The need for heavy post-processing of the data is due to the fact that the Higgs boson and other massive elementary particles are very short-lived. Typically they decay into other lighter particles in less than 10^{-21} seconds, i.e. long before they hit the detector. The signals observed in the detector originate from those decay products that are long-lived enough to reach the detector before further decays. These *final-state* particles are *reconstructed* and identified by finding correct combinations of detector signals.

The reconstruction of each collision event requires calibration and combination of signals from approximately 100 million readout channels of the CMS detector. The methods for particle reconstruction and identification are presented in Chapter 8.

Special emphasis is given to objects that are as demanding to reconstruct as they are critical for the $H^\pm$ search in hand: hadronically decaying tau leptons (τ_h), b jets, and missing transverse momentum, a variable that quantifies the momentum imbalance in the final state.

After storing and reconstructing the events, detailed simulations can be used to infer what type of processes possibly produced observed events. Involved statistical analysis is needed to test the compatibility of the data with the signal hypothesis (in our case, that events with $H^\pm$ are present in data) and with the background hypothesis (only SM events that resemble $H^\pm$ events are present). The statistical methods specifically designed for this purpose are presented in Chapter 9.

The search presented in this thesis is the first $H^\pm \to \tau^\pm \nu_\tau$ search with the CMS experiment at 13 TeV collision energy, based on the proton-proton collision data collected in 2016. The amount of data corresponds to an integrated luminosity of $35.9\,\mathrm{fb}^{-1}$. It targets the hadronic final state containing a τ_h, jets from hadronization of b quarks and light quarks, and missing transverse momentum due to neutrinos from the $H^\pm$ and τ decays.

The special feature of this final state is that the transverse mass of the $\tau\nu$ system can be reconstructed in an unambiguous way. If charged Higgs bosons exist, they are expected to show up as a Jacobian peak in this transverse mass distribution. The properties of this final state and the general analysis strategy are discussed further in Chapter 10.

To maximize the sensitivity of the analysis for a possible $H^\pm$ signal, the identification algorithms for different types of objects and the event selection criteria are re-optimized with respect to previous CMS analyses studying the same final state. A novel approach for categorization of collision events in the statistical analysis is introduced. The event selection and categorization criteria are detailed in Chapter 11.

As the goal of the analysis is to distinguish a possible small signal from an overwhelming background, the correct estimation of different background processes is critical. The background from Standard Model events can be split into two main components: events with a genuine τ_h, and events with a hadronic jet misidentified as τ_h. In Chapter 12, two alternative methods for estimation of the genuine-tau background are presented: one based on Monte Carlo simulation and another based on *tau embedding*, where events with muons are selected and the muons are replaced with simulated tau leptons.

The second background component, with jets misidentified as τ_h, cannot be reliably estimated from simulation, so a data-driven method that has also been further developed since the previous analyses, is presented. A third, small background contribution arises from events with electrons or muons misidentified as τ_h, and is estimated from simulation.

Each step in event reconstruction, identification of objects, and background estimation is subject to systematic effects that need to be recognized and corrected, and to systematic uncertainties that need to be included in the statistical analysis. All theoretical and experimental uncertainties taken into account in the analysis are described in Chapter 13.

Finally, the results of the search are presented in Chapter 14. For the first time in CMS the search range is extended up to $H^\pm$ mass of $3\,\text{TeV}$, and for the first time also the $H^\pm$ mass hypotheses close to the top quark mass are included in the search. The results are first presented in a model-independent way and then interpreted in the Minimal Supersymmetric Standard Model [34]. Chapter 15 summarizes the work and discusses the future prospects.

PART I

Theoretical background

Chapter 2

Standard Model

The Standard Model (SM) of particle physics describes the basic constituents of matter, the *elementary particles*, and their interactions via three fundamental forces: the electromagnetic, weak and strong forces. All physics in macroscopic scales is understood to emerge from elementary particles and their interactions, except for gravity which is not included in the SM but described by the theory of general relativity. However, in the energy regime relevant for the high energy physics experiments, the effect of gravity is negligible. The SM is formulated as a quantum field theory, combining the Glashow-Weinberg-Salam model of electroweak interactions [1–3] and quantum chromodynamics (QCD) of strong interactions [4–6]. The elementary particles obtain their masses via the so-called *Higgs mechanism* (Section 2.2), manifesting itself experimentally as the famous *Higgs boson*. In the following, a quick overview of the central concepts of the SM is provided. Comprehensive modern reviews of the SM can be found e.g. in Refs. [36, 37].

2.1 Particles and fields

The particle content of the SM is listed in Table 2.1. The measured masses of each particle, as well as internal properties of each particle such as spin, electric charge, and color charge, are listed. All matter is made of *quarks* and *leptons*, which come in three *generations* of different mass scales. The forces are mediated by *gauge bosons* (gluons, photons, $W^{\pm}$ and Z bosons). The gauge bosons are also called *vector bosons*, since these spin-1 particles correspond to vector fields in quantum field theory. The Higgs boson (H), on the other hand, is the only spin-0 *scalar* boson in the SM, corresponding to a scalar Higgs field described in Section 2.2.

Table 2.1: The particle content of the SM. Left-handed (L) quarks and leptons form $SU(2)_L$ doublets, whereas their right-handed (R) counterparts transform as singlets under $SU(2)$. The measured masses are listed based on Ref. [35]. In the SM, the neutrinos are assumed to be massless and exclusively left-handed. However, the 90% C.L. upper limit for the sum the three fermion masses ($m_{\nu,\text{tot}}$) is also included in the table. The $SU(3)_C$ representation governing the strong interaction, and the T_3, Y, and Q quantum numbers determining the electromagnetic and weak interactions are also listed.

	Particle(s)	Mass	Spin	$SU(3)_C$	T_3	Y	Q
Quarks	$\begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix}$	$m_u = 2.3\,\text{MeV}\ m_c = 1.2\,\text{GeV}\ m_t = 173\,\text{GeV}$ $m_d = 4.8\,\text{MeV}\ m_s = 0.1\,\text{GeV}\ m_b = 4.7\,\text{GeV}$	$\frac{1}{2}$	3	$\frac{1}{2}$ $-\frac{1}{2}$	$\frac{1}{6}$ $\frac{1}{6}$	$\frac{2}{3}$ $-\frac{1}{3}$
	$u_R,\ c_R,\ t_R$				0	$\frac{2}{3}$	$\frac{2}{3}$
	$d_R,\ s_R,\ b_R$				0	$-\frac{1}{3}$	$-\frac{1}{3}$
Leptons	$\begin{pmatrix} \nu_{e,L} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu,L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau,L} \\ \tau_L \end{pmatrix}$	$m_{\nu,\text{tot}} \lesssim 0.2\,\text{eV}$ $m_e = 0.51\,\text{MeV}\ m_\mu = 0.11\,\text{GeV}\ m_\tau = 1.78\,\text{GeV}$	$\frac{1}{2}$	1	$\frac{1}{2}$ $-\frac{1}{2}$	$-\frac{1}{2}$ $-\frac{1}{2}$	0 -1
	$e_R,\ \mu_R,\ \tau_R$				0	-1	-1
Gauge bosons	g	$m_g = 0$	1	8	0	0	0
	$W^\pm$	$m_W \approx 80\,\text{GeV}$	1	1	± 1	0	± 1
	Z	$m_Z \approx 91\,\text{GeV}$	1	1	0	0	1
	γ	$m_\gamma = 0$	1	1	0	0	0
Higgs boson	h	$m_h \approx 125\,\text{GeV}$	0	1	$-\frac{1}{2}$	$\frac{1}{2}$	0

In quantum field theory, the quantum fields and their interactions included in a theory can be summarized by defining the *Lagrangian density* (simply "Lagrangian" from hereon), from which the particle masses and interactions can be derived in detail. Particles arise as excitations of the corresponding fields. The SM Lagrangian consists of five parts:

$$\mathcal{L} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{quark}} + \mathcal{L}_{\text{lepton}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}, \tag{2.1}$$

which describe the gauge, quark, lepton, Higgs and Yukawa *sectors* of the theory. In the following, each of these sectors will be introduced, and their phenomenological consequences outlined.

The SM Lagrangian is invariant under the symmetry group $U(1)_Y \times SU(2)_L \times SU(3)_C$. The first two symmetries govern the electroweak interactions. The generator of the $U(1)_Y$ symmetry is the *weak hypercharge Y* and the corresponding gauge field is denoted as B_μ. The three generators T_i of $SU(2)_L$ are the three components of *weak isospin T*, and the corresponding gauge three fields are W_μ^α ($\alpha = 1, 2, 3$). The weak hypercharge

and the third component of the weak isospin define the electric charge Q of a particle:

$$Q = Y + T_3. \tag{2.2}$$

The weak interactions are found to discriminate between particles based on their *chirality*, a property of a fermion that defines the representation of the Lorentz group under which its corresponding Dirac spinor field transforms [37]. Since the weak interactions affect only "left-handed" fermions, i.e. those with negative chirality, the L here denotes "left". It reminds us that only left-handed fermions form $SU(2)_L$ doublets, while the right-handed fermions are $SU(2)$ singlets.

Finally, QCD is invoked by the eight generators of the $SU(3)_C$, corresponding to different *color charges*, with eight gauge fields denoted as G_μ. We denote the coupling strengths corresponding to $U(1)_Y$, $SU(2)_L$, and $SU(3)_C$, as g_1, g_2 and g_3, respectively.

The kinetic terms for all these gauge fields, respecting the gauge symmetries, can be written as

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{8} \text{Tr} W_{\mu\nu} W^{\mu\nu} - \frac{1}{2} \text{Tr} G_{\mu\nu} G^{\mu\nu}, \tag{2.3}$$

where the *field strength tensors* are defined as

$$\begin{aligned}
B_{\mu\nu} &= \partial_\mu B_\nu - \partial_\nu B_\mu, \\
W_{\mu\nu} &= \partial_\mu W_\nu - \partial_\nu W_\mu + ig_2(W_\mu W_\nu - W_\nu W_\mu), \\
G_{\mu\nu} &= \partial_\mu G_\nu - \partial_\nu G_\mu + ig_3(G_\mu G_\nu - G_\nu G_\mu).
\end{aligned} \tag{2.4}$$

The quarks come in six *flavors*. The three up-type quarks, i.e. up (u), charm (c) and top (t) quarks carry an electric charge of $+1/3$, while the three down-type quarks, i.e. down (d), strange (s), and bottom (b) quarks have a charge of $-2/3$. The interactions of quarks with other particles are dictated by the second term of Eq. (2.1)

$$\mathcal{L}_{\text{quark}} = (\bar{u}_L, \bar{d}_L)\tilde{\sigma}^\mu i D_\mu \begin{pmatrix} u_L \\ d_L \end{pmatrix} + \bar{u}_R \sigma^\mu i D_\mu u_R + \bar{d}_R + \sigma^\mu i D_\mu d_R + \text{h.c.} \tag{2.5}$$

where the Pauli matrices σ^α ($\alpha = 1, 2, 3$) define $\tilde{\sigma}^\mu = (1, -\sigma^1, -\sigma^2, -\sigma^3)$ and $\sigma^\mu = (1, \sigma^1, \sigma^2, \sigma^3)$, and h.c. means Hermitian conjugate. Here we implicitly sum over the three quark generations for quarks, so for example u_L refers to any left-handed

up-type quark. The covariant derivatives are defined such that $\mathcal{L}_{\text{quark}}$ respects the gauge symmetries, and they reveal the couplings of quarks to gauge fields:

$$
\begin{aligned}
D_\mu \begin{pmatrix} u_L \\ d_L \end{pmatrix} &= (\partial_\mu + \frac{ig_1}{6}B_\mu + \frac{ig_2}{2}W_\mu + ig_3 G_\mu) \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \\
D_\mu u_R &= (\partial_\mu + \frac{i2g_1}{3}B_\mu + ig_3 G_\mu)u_R, \\
D_\mu d_R &= (\partial_\mu + \frac{-ig_1}{3}B_\mu + ig_3 G_\mu)d_R.
\end{aligned}
\tag{2.6}
$$

We observe that all quarks couple to B_μ and G_μ, and hence they are subject to the electromagnetic and strong interactions. Only left-handed quarks couple to W_μ, so right-handed quarks are not included in weak interactions.

The leptons also come in six *flavors*. They have integer electric charge and no color charge. The three charged leptons (electron, muon and tau lepton) have all electric charge of -1 but they differ in mass. They are accompanied by corresponding three neutral leptons called *neutrinos*: electron neutrino (ν_e), muon neutrino (ν_μ) and tau neutrino (ν_τ). The neutrinos have masses very close to zero, and they interact only via the weak interactions. In the SM, all neutrinos are left-handed and massless, whereas other leptons and quarks are massive and exist in both left-handed and right-handed versions. The leptonic part of the SM Lagrangian is

$$
\mathcal{L}_{\text{lepton}} = (\bar{\nu}_L, \bar{e}_L)\tilde{\sigma}^\mu iD_\mu \begin{pmatrix} \nu_L \\ e_L \end{pmatrix} + \bar{e}_R \sigma^\mu iD_\mu e_R + \bar{\nu}_R \sigma^\mu iD_\mu \nu_R + \text{h.c.}
\tag{2.7}
$$

where the interactions with gauge bosons enter again via the covariant derivatives

$$
D_\mu \begin{pmatrix} \nu_L \\ e_L \end{pmatrix} = (\partial_\mu + \frac{ig_1}{2}B_\mu + \frac{ig_2}{2}W_\mu) \begin{pmatrix} \nu_L \\ e_L \end{pmatrix},
\tag{2.8}
$$

$$
D_\mu e_R = (\partial_\mu + igB_\mu)e_R.
\tag{2.9}
$$

Since leptons are SU(3) singlets, they do not experience strong interactions, so G_μ is not included in the covariant derivative. As previously, only left-handed particles are subject to the weak interaction represented by W_μ.

For quarks (leptons), each generation corresponds to a pair of up-type and down-type quark (charged lepton and its neutrino). The particles of second and third generation

are short-lived, since they decay into lighter particles. Only the first generation particles are stable, since they cannot decay into lighter particles.

All quarks and leptons are *fermions*, as they have spin of 1/2, and hence follow the Pauli exclusion principle: two fermions cannot occupy the same quantum state simultaneously [38]. Particles with integer spin, not subject to the exclusion principle are called *bosons*. Each fermion also has its own antiparticle, with opposite electric charge but the same mass and spin. In this thesis, charge conjugation is implied, so for example by "top quark" we refer to both top quark and antiquark.

The three fundamental forces – electromagnetic, weak, and strong interactions – are mediated by gauge bosons (vector bosons), which all have spin 1. The massless *gluons* mediate the strong interaction, and the massless *photons* mediate the electromagnetic interaction. The weak interactions, on the other hand, are mediated by massive vector bosons $W^\pm$ and Z. Before taking a closer look at these different interactions, let us discuss how the $W^\pm$ and Z obtain their masses.

2.2 Electroweak symmetry breaking

Naively, one would expect the four bosons mediating electroweak interactions ($W^\pm$, Z, γ) to correspond to the four fields involved in the electroweak interactions: B_μ and W_μ^α ($\alpha = 1, 2, 3$). However, the requirement of local gauge invariance forces all the four bosons corresponding to these fields to be massless, since any mass term of type $m^2 W_\mu W^\mu$ would break the gauge invariance. Still, experimentally the $W^\pm$ and Z bosons are known to be massive: $m_{W^\pm} \approx 80.4\,\text{GeV}$ and $m_Z \approx 91.2\,\text{GeV}$ [35].

The dilemma is solved by the Englert-Brout-Higgs-Guralnik-Hagen-Kibble mechanism [7–12] referred to simply as the Higgs mechanism from hereon. This mechanism allows introducing the mass term in an indirect way via spontaneous symmetry breaking. The idea is to add a single complex scalar doublet, called *Higgs doublet*,

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1(x) + i\phi_2(x) \\ \phi_3(x) + i\phi_4(x) \end{pmatrix}, \tag{2.10}$$

with four real components $\phi_1, \ldots, \phi_4$. The Lagrangian for this field contains a kinetic term and a potential V:

$$\mathcal{L}_{\text{Higgs}} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi), \tag{2.11}$$

Figure 2.1: "Mexican hat" shaped Higgs potential $V(\Phi)$ shown as a function of two real parameters ϕ_3 and ϕ_4. The usual choice of minimum is illustrated with a blue ball.

where the covariant derivative

$$D_\mu \Phi = \partial_\mu + (\frac{ig_1}{2} Y B_\mu + \frac{ig_2}{2} W_\mu^\alpha \sigma^\alpha) \Phi \tag{2.12}$$

preserves the local gauge invariance of $\mathcal{L}$ and couples the Higgs doublet to B_μ and W_μ via coupling constants g_1 and g_2, respectively. The most general renormalizable and gauge invariant potential is

$$V(\Phi) = -\mu^2(\Phi^\dagger \Phi) + \lambda(\Phi^\dagger \Phi)^2, \tag{2.13}$$

where μ^2 has a dimension of mass squared and λ is a dimensionless real parameter. To keep the vacuum stable, the potential needs to be bounded from below, so $\lambda > 0$. For $\mu^2 < 0$, the potential has a unique minimum and it preserves the gauge symmetry of the electroweak Lagrangian. For $\mu^2 > 0$, the potential has a continuum of non-zero minima. This famous "Mexican hat" potential is illustrated in Figure 2.1. The minima of the potential are found by requiring $\partial V/\partial \Phi = \partial V/\partial \Phi^\dagger = 0$ and $\partial^2 V^2/(\partial \Phi \partial \Phi^\dagger) = 0$ in a minimum $\langle \Phi \rangle$, yielding

$$|\langle \Phi \rangle|^2 = \frac{1}{2}(\langle \phi_1 \rangle^2 + \langle \phi_2 \rangle^2 + \langle \phi_3 \rangle^2 + \langle \phi_4 \rangle^2) = \frac{\mu^2}{2\lambda} \equiv \frac{v^2}{2}. \tag{2.14}$$

The real fields $\phi_1, \ldots, \phi_4$ can be chosen freely as long as the above condition is satisfied. A standard choice is $\langle\phi_1\rangle = \langle\phi_2\rangle = \langle\phi_4\rangle = 0$, $\langle\phi_3\rangle^2 = v^2$, which gives

$$\langle\Phi\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \tag{2.15}$$

This is the *vacuum expectation value* of the Higgs doublet. The choice of this particular minimum leads to the spontaneous breaking of the local $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ gauge symmetry to a residual $\mathrm{U}(1)_{\mathrm{EM}}$ symmetry. Now we can rewrite the Lagrangian in Eq. (2.11) by expanding $\Phi(x)$ around its minimum as:

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \tag{2.16}$$

where $h(x)$ is called the *Higgs field*. The expanded Lagrangian becomes

$$\begin{aligned}
\mathcal{L}_{\mathrm{Higgs}} =& \frac{1}{2}(\delta_\mu h)(\delta^\mu h) + \frac{1}{2}\mu^2 h^2 \\
&+ \frac{g_2^2 v^2}{8}(W_{1,\mu}W_1^\mu + W_{2,\mu}W_2^\mu) \\
&+ \frac{v^2}{8}(g_1 B_\mu - g_2 W_{3,\mu})(g_1 B^\mu - g_2 W_3^\mu) + \text{interaction terms}.
\end{aligned} \tag{2.17}$$

We identify the first two terms as the kinetic term and the mass term of a new particle, the Higgs boson. The mass squared of the Higgs boson is $m_{\mathrm{h}}^2 = 2\mu^2 = 2\lambda v^2$. To interpret the remaining terms, we need to perform transformations that reveal the mass eigenstates corresponding to physical particles, which are linear combinations of the original gauge fields. Starting with the third term, the mass eigenstates are found to be linear combinations of $W_{1,\mu}$ and $W_{2,\mu}$, which we denote as

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_{1,\mu} \mp iW_{2,\mu}). \tag{2.18}$$

In terms of $W_\mu^\pm$ fields the third term in Eq. (2.17) is

$$(\frac{g_2 v}{2})^2 (W^+)_\mu (W^-)^\mu, \tag{2.19}$$

which we identify as a mass term for charged gauge bosons $W^\pm$ with mass $m_{\mathrm{W}^\pm} = g_2 v/2$. In the fourth term, the hypercharge gauge field $W_{3,\mu}$ and the isospin field B_μ get mixed.

The mass eigenstates result from a transformation that corresponds to a unitary rotation by the *weak mixing angle* (the Weinberg angle) θ_W:

$$\begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} B_\mu \\ W_{3,\mu} \end{pmatrix}. \tag{2.20}$$

After the rotation, the fourth term in Eq. (2.17) becomes

$$\frac{1}{2}\left(\frac{vg_2}{2\cos\theta_W}\right)^2 Z_\mu Z^\mu + 0 \times A_\mu A^\mu, \tag{2.21}$$

where the first part corresponds to a mass term for a neutral gauge boson, the Z boson, with a mass of $m_Z = vg_2/(2\cos\theta_W)$, and the second part to a massless gauge boson, i.e. the photon. Consequently, the $W^\pm$ and Z boson masses are related by the weak mixing angle:

$$\cos\theta_W = \frac{m_{W^\pm}}{m_Z}. \tag{2.22}$$

Hence by measuring m_Z and $m_{W^\pm}$, one can infer θ_W, measured to be $\approx 28°$, and the value of v, measured to be $\approx 246\,\text{GeV}$.

To summarize, the Higgs mechanism can explain the masses of the $W^\pm$ and Z bosons, at the cost of spontaneously breaking the electroweak symmetry. It also predicts the existence of a scalar boson, but its mass is a free parameter of the theory (since μ^2 is a free parameter in Eq. (2.13)) that needs to be measured experimentally.

The four degrees of freedom of the complex doublet Φ were transformed into one degree of freedom for the new scalar field h, and into three new longitudinal degrees of freedom absorbed by the massive gauge bosons $W^\pm$ and Z. The photon remains massless due to a residual $U(1)_{EM}$ symmetry, so it has only transverse degrees of freedom.

In addition, the Higgs field can generate masses for all fermions. Since the left-handed fermions (f_L) form $SU(2)_L$ doublets, while the right handed fermions (f_R) are $SU(2)_L$ singlets, a naive fermion mass term of type $m(\bar{f}_L f_R + \bar{f}_R f_L)$ is not allowed by the gauge symmetry. However, *Yukawa couplings* of the form

$$\mathcal{L}_{\text{Yukawa}} = -y_f\left((\bar{f}_L\Phi)f_R + \text{h.c.}\right) \tag{2.23}$$

that also contain the Higgs field Φ, are allowed. Here y_f is the Yukawa coupling constant of a fermion f. Expanding around $\langle \Phi \rangle$ using Eq. (2.16), we find

$$
\begin{aligned}
\mathcal{L}_{\text{Yukawa}} &= -\frac{y_f}{\sqrt{2}} v \left(\bar{f}_L f_R + \bar{f}_R f_L \right) - \frac{y_f}{\sqrt{2}} \left(\bar{f}_L f_R + \bar{f}_R f_L \right) h(x) \\
&= -m_f \left(\bar{f}_L f_R + \bar{f}_R f_L \right) - \frac{m_f}{v} \left(\bar{f}_L f_R + \bar{f}_R f_L \right) h(x).
\end{aligned}
\tag{2.24}
$$

In the second line, the mass of fermion f is identified to be $m_f = v y_f / \sqrt{2}$. Hence the mass of a fermion is proportional to the corresponding Yukawa coupling. As the Yukawa coupling constants y_f appear as new free parameters, the Higgs mechanism (and the SM) does not predict the fermion masses or any relations between them. The second term describes the coupling of the Higgs field h to fermions, which is proportional to the mass of the fermion. Couplings proportional to the fermion mass are a distinctive feature of the Higgs boson. For example, the massive gauge bosons couple to all leptons with identical strength — this property of the SM is known as *lepton universality*.

In 2012, the ATLAS and CMS experiments observed a Higgs boson with a mass of approximately 125 GeV [14–16], providing strong evidence for the Higgs mechanism. The discovery was followed by precision measurements of the mass, couplings, spin, parity, and production rates of the observed boson, which were found to be consistent with the predictions of the SM [39–43].

2.3 Electromagnetic and weak interactions

When the rotation of B_μ and $W_{3,\mu}$ into A_μ and Z_μ by angle θ_W, as defined in Eq. (2.20), is applied to the interaction term governing the electroweak force, it decomposes into separate terms corresponding to electromagnetic interactions mediated by massless *photons*, and neutral weak currents mediated by massive Z bosons. The charged weak currents are mediated by massive $W^\pm$ bosons.

In other words, after the spontaneous symmetry breaking, the electromagnetic and weak interactions appear as separate interactions, even though they have a common origin. The decomposition also connects the coupling strengths g_1 and g_2 to each other and to the elementary electric charge e:

$$
g_1 \sin \theta_W = g_2 \cos \theta_W = e.
\tag{2.25}
$$

Since photons are massless, the electromagnetic force has an infinite range. All particles carrying non-zero electric charge, i.e. quarks, charged leptons, and $W^{\pm}$ bosons, are subject to electromagnetic interactions.

As discussed above, the weak interaction only affects "left-handed" fermion doublets with negative chirality. For quarks, the $W^{\pm}$ bosons do not couple directly to the doublets made of mass eigenstates of (left-handed) up-type and down-type quarks. Instead, they couple to doublets of *flavor eigenstates*, which are linear combinations of the mass eigenstates as specified by the Cabibbo–Kobayashi–Maskawa (CKM) matrix V_{CKM}. The CKM matrix is a 3×3 unitary matrix which connects the three generations of quark mass eigenstates to three generations of flavor eigenstates. It has three real parameters and one complex phase.

The phase in the CKM matrix implies that weak interactions violate the CP symmetry, which is the product of the C-symmetry (charge conjugation symmetry) and P-symmetry (parity symmetry, i.e. symmetry under the flip in the sign of one spatial coordinate). The nonzero off-diagonal values of the CKM matrix correspond to *flavour-changing charged currents* mediated by the $W^{\pm}$ bosons. On the other hand, it can be shown that the Z bosons and photons cannot mediate *flavor-changing neutral currents* at tree level, and even in loop processes they are suppressed by the so-called Glashow–Iliopoulos–Maiani mechanism [44].

As discussed above, the neutrinos interact only via weak interaction. Since in the SM they are assumed to be massless, the mass eigenstates and flavor eigenstates of the three left-handed charged lepton doublets coincide. However, we know experimentally that neutrinos have non-zero masses, and their mass eigenstates differ from the flavor eigenstates. This mixing in the lepton sector is parameterized by the Pontecorvo–Maki–Nakagawa–Sakata matrix, which is analogous to the CKM matrix. Since the neutrino masses are outside the scope of the SM, we return to them in Chapter 3.

2.4 Strong interactions

As the structure of the SU(3) symmetry group that governs the strong interactions can be understood via an analogy with the RGB (red/blue/green) color system, the theory of strong interactions is called *quantum chromodynamics* (QCD). The strong force acts on *color charge* carried by quarks and gluons, and the color charge is conveniently expressed in terms of red, green and blue "colors" of quarks and corresponding "anti-colors" of antiquarks.

A local gauge symmetry always introduces a number of massless bosons, corresponding to the number of generators of the group. In case of the $SU(3)$ group of QCD, its eight generators correspond to eight *gluons* that mediate the strong interactions.

Gluons are massless, but unlike photons with zero electric charge, or the $W^\pm$ and Z bosons with zero weak hypercharge, they carry color charge themselves. This gives rise to gluon self-interactions, which polarize the vacuum so that the strength of the strong force increases with distance, as opposed to the electromagnetic or weak forces which become weaker as distance increases. This means that quarks and gluons become (almost) free particles only at very large energy scales. This property of the QCD is referred to as *asymptotic freedom*.

The gluon self-interactions give rise to rich phenomenology in hadron colliders, which is very different from the phenomenology of electromagnetic and weak interactions. Due to asymptotic freedom, at limited energy scales we never observe quarks as free particles. Instead, they appear as color neutral bound states (hadrons), typically consisting of a quark and antiquark (mesons) or three quarks (baryons). Therefore in high-energy physics experiments, instead separate of quarks or gluons, collimated sprays of hadrons called *jets* are observed.

The total momentum of a hadron is not shared only between the *valence quarks* which define the type of hadron, but also with the *sea quarks* that are constantly created and annihilated inside the hadron, as well as by gluons exchanged among all quarks inside the hadron.

CP-violation has not been observed in strong interactions, even though it is not forbidden by the structure of the theory. Therefore CP symmetry in QCD seems to be an accidental symmetry. Its origin is an open question, known as *the strong CP-problem.*

2.5 Properties of tau leptons

The distinctive feature of the tau lepton is its relatively large mass of $m_\tau = 1.777\,\text{GeV}$ [35], which is approximately 17 times larger than the mass of the muon and 3500 times larger than the mass of the electron. This means that the coupling of Higgs bosons to tau leptons is much stronger than to other leptons. Therefore Higgs boson decay channels involving tau leptons can have significant branching fractions and they are of considerable experimental interest. This also applies to the $H^\pm \to \tau^\pm \nu_\tau$ channel.

The large mass also implies a short lifetime, measured to be $\tau_\tau = 0.29\,\text{ps}$, corresponding to $c\tau_\tau = 87\,\mu\text{m}$. For very energetic tau leptons moving close to the speed of light, the mean distance traveled in the laboratory frame differs from $c\tau_\tau$ by the Lorentz factor $\gamma = E/m_\tau$ due to relativistic time dilation. For example, for a tau lepton with an energy of $E = 50\,\text{GeV}$, the mean flight distance in the detector is $\approx 2{,}5\,\text{mm}$.

Tau leptons are heavy enough to decay both into lighter leptons and into hadrons. In the following, we discuss these different decay modes, as well as the helicity correlations which can provide information about the origin of the tau lepton, providing a handle to discriminate between tau leptons coming from $W^\pm$ and $H^\pm$ bosons.

2.5.1 Decay modes

The most common final states of tau lepton decays and their *branching fractions* are listed in Table 2.2. Here $h^\pm$ refers to a hadron, either $\pi^\pm$ or $K^\pm$. All decay modes are characterized by the presence of a tau neutrino. As discussed later in Chapter 8, neutrinos cannot be detected directly in collider experiments, but their presence can be inferred indirectly since they affect the momentum balance in the event.

For the leptonic τ decay modes, the momentum imbalance is the only indicator about the origin of the electrons and neutrinos, since the properties of these particles alone are similar to prompt electrons and muons coming directly from $W^\pm$ or Z decays. The hadronic decays proceed via an intermediate resonance, such as ρ (720 MeV) or a_1 (1260 MeV), or directly into final states that include charged pions or kaons and possibly also neutral pions. These particles form the visible part of the τ decay products, which is a collimated jet of hadrons denoted as τ_h. Advanced identification algorithms are needed to experimentally distinguish such τ_h signatures from quark and gluon jets.

The hadronic tau lepton decays are classified according to the number of *prongs*, i.e. the number of charged hadrons in the decay mode. The analysis presented in this thesis targets the one-prong hadronic decays, which constitute 48% of all tau lepton decays and 75% of all hadronic decays.

2.5.2 Helicity correlations

The *helicity* of a particle is defined as the sign of the projection of its spin onto the direction of its momentum. If the spin projection is parallel (antiparallel) to the

Table 2.2: Branching fractions ($\mathcal{B}$) for different tau lepton decay channels, based on Ref. [35]. The uncertainties for all values are < 0.1 percent unit, except for the last one which is determined by requiring that the branching fractions sum to unity. Charge conjugation is implied.

τ decay channel	Resonance	$\mathcal{B}$ (%)	Total (%)
leptonic decay modes			35.6
$\tau^- \to e^- \bar{\nu}_e \nu_\tau$		17.8	
$\tau^- \to \mu^- \bar{\nu}_\mu \nu_\tau$		17.4	
hadronic decay modes			
one-prong decays (excl.K^0)			48.2
$\tau^- \to h^- \nu_\tau$		11.5	
$\tau^- \to h^- \pi^0 \nu_\tau$	ρ^-	25.9	
$\tau^- \to h^- 2\pi^0 \nu_\tau$	a_1^-	9.6	
$\tau^- \to h^- \nu_\tau + \geq 3\pi^0$		1.3	
three-prong decays (excl.K^0)			14.6
$\tau^- \to h^- h^- h^+ \nu_\tau$	a_1^-	9.5	
$\tau^- \to h^- h^- h^+ \nu_\tau + \geq 1\pi^0$		5.1	
five-prong decays (excl.K^0)			0.1
$\tau^- \to 5h^- 2h^+ \nu_\tau + \geq 0\pi^0$		0.1	
others (incl. decays with K^0)			≈ 2

momentum vector, the particle is said to be right-handed (left-handed) and associated with a *polarization* of $+1$ (-1). Unlike chirality, the helicity of a particle depends on the chosen reference frame, but for highly relativistic particles ($E \gg m$) the two agree.

The helicity of tau leptons carries information about their origin, so it provides a useful handle for discrimination between the tau leptons originating from $H^\pm$ decays and those originating from $W^\pm$ decays. Since the charged Higgs boson is a scalar particle, and neutrinos are always left-handed (in the SM), the tau lepton produced in $H^\pm$ decay is always right-handed (polarization $P_\tau = +1$). On the contrary, when the vector boson $W^\pm$ decays into $\tau^\pm \nu_\tau$, the fact that the neutrino is left-handed forces also the tau lepton to be left-handed ($P_\tau = -1$).

In the following we concentrate on 1-prong hadronic tau decays. We argue how the origin of the τ lepton affects different decay modes, and motivate a simple kinematic

selection that can be used to enhance the $H^\pm \to \tau^\pm \nu_\tau$ signal. The helicity correlations and selection techniques based on them are discussed in more detail in Ref. [45–48].

Let us concentrate on three decay channels: $\tau^- \to \pi^- \nu_\tau$, $\tau^- \to \rho^- \nu_\tau \to \pi^- \pi^0 \nu_\tau$, and $\tau^- \to a_1^- \nu_\tau \to \pi^- 2\pi^0 \nu_\tau$, which constitute $> 95\%$ of all one-prong τ decays, and neglect the decay modes with $\geq 3\pi^0$'s.

In the $\tau^- \to \pi^- \nu_\tau$, the angle θ between the pion and the τ spin quantization axis, defined to coincide with the τ direction in the laboratory frame, is distributed as

$$\frac{1}{\Gamma_\pi} \frac{d\Gamma_\pi}{d\cos\theta} = \frac{1}{2}(1 + P_\tau \cos\theta), \tag{2.26}$$

where the differential decay width $d\Gamma_\pi / d\cos\theta$ is measured in the rest frame of the τ. As P_τ is $+1$ (-1) for a tau lepton originating from a $H^\pm$ $(W^\pm)$, a forward (backward) emission of the pion is favored in the τ rest frame. In the laboratory frame, this implies that the pions originating from $H^\pm$ have on average higher momentum than those originating from $W^\pm$.

The two other decay channels proceed via intermediate resonances ρ (pseudovector meson) or a_1 (vector meson), which we commonly denote as v. In case of longitudinal polarization (L), the angular distribution is

$$\frac{1}{\Gamma_{v,L}} \frac{d\Gamma_\pi}{d\cos\theta} = \frac{1}{2}\frac{m_\tau^2}{m_\tau^2 + 2m_v^2}(1 + P_\tau \cos\theta), \tag{2.27}$$

while for the transversely-polarized (T) case it is

$$\frac{1}{\Gamma_{v,T}} \frac{d\Gamma_\pi}{d\cos\theta} = \frac{m_v^2}{m_\tau^2 + 2m_v^2}(1 - P_\tau \cos\theta). \tag{2.28}$$

Therefore for the longitudinally polarized mesons, the behavior is similar as above: the ones originating from $H^\pm$ are "harder", i.e. they tend to have higher momentum. For the transversely polarized mesons, the opposite is true: the ones originating from $W^\pm$ are harder. Hence a variable that discriminates between transversely and longitudinally polarized τ candidates can be used for discrimination between τ candidates from $H^\pm$ and $W^\pm$. The distribution of the angle θ in Eqs. (2.26)–(2.28) can be expressed as

$$\cos\theta = \frac{2X - 1 - \frac{m'^2}{m_\tau^2}}{1 - \frac{m'^2}{m_\tau^2}}, \tag{2.29}$$

Figure 2.2: Distributions of the τ decay widths as a function of the momentum fraction X for $\tau^- \to \rho^- \nu_\tau \to \pi^- \pi^0 \nu_\tau$ and $\tau^- \to a_1^- \nu_\tau \to \pi^- 2\pi^0 \nu_\tau$ decay channels, shown for cases where the intermediate resonance is longitudinally (L) or transversely (T) polarized, calculated in the τ rest frame. [47]

where X is the fraction of the τ_h momentum carried by the charged pion that characterizes the one-prong decays: $X = p_{\pi^\pm}/p_{\tau_h}$, defined in the τ rest frame, and m' is the mass of the charged pion (resonance meson) in the case of Eq. (2.26) (2.27 and 2.28). The distribution of X for both polarization modes of ρ and a_1 mesons is shown in Figure 2.2. We observe that in the case of longitudinal polarization, the charged pion carries either a large (> 0.7) or a small (< 0.3) fraction of the τ momentum, while transverse polarization the sharing of momentum is typically more even.

For the $\pi^- \nu_\tau$ decay mode, the distribution would peak sharply at 1 since in this mode all (visible) τ_h momentum is carried by the charged pion. We notice that by selecting only events with e.g. $X > 0.75$, we can select a relatively pure sample of $\rho_T \nu_\tau$ and $\pi^- \nu_\tau$ decays. If a suitable momentum selection has been applied, most of these tau leptons originate from $H^\pm$, while most of the transversely-polarized or lower-momentum tau leptons from $W^\pm$ decays are rejected. It is worth noting, however, that the majority of the left-handed $a_{1L} \nu$ decays that also (mostly) originate from $H^\pm$, are sacrificed in the process.

In practice, the variable X is usually defined in the laboratory frame instead of the τ rest frame, in which case it is often called R_τ. Simulation studies that go beyond the simple argument presented here, including the nonresonant contribution to hadronic τ decays and realistic signal and background simulations, have confirmed that such a selection indeed improves the sensitivity for $H^\pm$ signal [48]. In addition to suppressing genuine τ leptons coming from $W^\pm$, this type of selection also reduces the background contribution from quark or gluon jets misidentified as τ_h. In the second part of this thesis, when a selection based on R_τ is applied in the data analysis, its power will be demonstrated.

Chapter 3

Beyond the Standard Model

The efforts to formulate theories beyond the Standard Model (BSM) are motivated both by certain theoretically unsatisfactory aspects of it, and by empirical observations that cannot be explained in the framework of the SM. Before discussing some possible extensions of the SM, let us briefly review some of its shortcomings. A more thorough overview of the subject is provided e.g. in Ref. [49].

One of the main theoretical puzzles related to the SM is the *hierarchy problem*. The higher-order quantum corrections to the Higgs boson mass-squared parameter μ^2 involve quadratic divergences. These divergences make the Higgs potential sensitive to physics in arbitrarily large energy scales, including the Planck scale Λ_P where the gravity enters. Even though gravity is not included in the SM, and the details of Planck-scale quantum physics are not yet known, the fact that the Higgs potential is sensitive to Λ_P scale phenomena means that extreme fine-tuning of the order of

$$\frac{\Lambda_P^2}{v^2} = \frac{(10^{19}\,\text{GeV})^2}{(10^2\,\text{GeV})^2} = 10^{34} \tag{3.1}$$

would be needed in the parameters describing the Planck scale physics to obtain the measured (low) value of $v^2 \approx (246\,\text{GeV})^2$. An extended theory could provide a mechanism that naturally keeps v^2 in the observed scale, without excessive fine-tuning.

Another interesting question is the stability of the electroweak vacuum. The measured masses of the SM parameters, especially those of the Higgs boson and the top quark, suggest that the observed electroweak vacuum v is *metastable*, i.e. that it represents a local rather than global energy minimum. This metastable state has a finite lifetime, and there is a non-zero probability that quantum tunneling to the global minimum occurs. Therefore a stable vacuum is considered theoretically preferable. Introducing

physics beyond the SM, however, can radically change the analysis of the situation and make the observations consistent with a stable vacuum.

The SM contains 19 free parameters that seem arbitrary, as their values are not explained or guided by the theory. All of these parameters, however, do not look arbitrary, and the observations suggest that there are symmetries in nature which are not contained in the SM. For example, the fermion masses arrange the fermions into three generations which have a definite mass hierarchy not explained by the SM. Another example is the *strong CP problem*: quantum chromodynamics (QCD) is experimentally found to preserve the CP-symmetry, although there is no fundamental symmetry in the SM that would require this.

At low energy scales, the SM describes three distinct types of interactions. The Glashow-Weinberg-Salam theory manages to unify the electromagnetic and weak interactions, suggesting that their coupling constants become comparable at high energies. The strong interaction, on the other hand, is completely separate from the other two. While the coupling constants of the electromagnetic and weak interactions are related by Eq. (2.25), the strong coupling constant is an independent parameter and it does not converge with the other two even at high energies. An extended model could provide unification of all three forces.

Gravitation is not included in the SM, since attempts to build a predictive quantum theory of gravity (either unified with the other forces or not) have not been successful. At the same time, the classical description of gravity by the general relativity involves a mystery too: in the framework of general relativity, the observed accelerating expansion of the universe is accommodated by choosing a specific value for a parameter called *cosmological constant*. This parameter is understood to represent *dark energy*, a specific type of energy that is present everywhere at a uniform energy density and causes a negative pressure that dilutes matter. The exact form and origin of this energy is, however, not understood.

In 1933 it was observed that the rotation velocities of certain galaxies do not match the amount of luminous matter they are observed to contain, indicating that they contain *dark matter* that does not interact electromagnetically and thus escapes optical detection. These days the existence of dark matter is an established fact, confirmed by a large amount of astrophysical and cosmological evidence, including measurements of gravitational lenses, baryon acoustic oscillations and the cosmic microwave background (CMB) [35]. The current standard model of cosmology suggests that dark

matter constitutes 26% of the mass and energy in the universe. 69% is attributed to the dark energy discussed above, and only 5% to the ordinary matter described by the SM [35].

As the ordinary matter around us is almost exclusively matter and not antimatter, it is natural to ask for a physical explanation for this imbalance. The CMB observations suggest that matter dominates over antimatter everywhere in the observable universe. In 1967, Andrei Sakharov formulated three conditions for baryogenesis where matter and antimatter are produced at different rates [50]. These conditions include violation of CP-symmetry. Although the SM introduces CP-violation via the complex phase of the CKM matrix, its magnitude is not large enough to explain the observed matter-antimatter-asymmetry [35].

The SM contains three left-handed neutrinos, which are massless and correspond to the three different lepton flavors. The observation of *neutrino oscillations* however suggests that neutrinos have non-zero mass, and their mass eigenstates differ from the flavor eigenstates. The mass differences between different neutrino mass eigenstates have been measured, and upper limits on their masses have been set [35]. While the neutrino oscillations can be plugged in the SM at the cost of introducing a number of additional free parameters (masses and mixing angles in analogy to the quark mixing via the CKM matrix), the nature of the neutrino mass remains an open question.

These and other shortcomings of the SM suggest that it is not a complete theory, but a low-energy limit of a more general theory. Regarding the form of the more general theory, a large number of hypotheses have been presented over decades. As most of the extended models include a more complex Higgs sector than the one in the SM, we continue by discussing possible ways of extending the SM Higgs sector.

3.1 Models with extended Higgs sectors

The Higgs sector of the SM is *minimal*, in the sense that it contains only one complex scalar doublet, introducing a minimal amount of new fields and corresponding free parameters needed for the electroweak symmetry breaking. However, from a theoretical point of view, the minimal Higgs sector is an arbitrary choice. Over history, the experiments have revealed that the other sectors in the SM are far from minimal versions, provoking reactions like the famous "who ordered that?" exclamation by Nobel laureate Isidor Isaac Rabi after the discovery of the muon in 1936.

Therefore models with extended Higgs sectors, containing more fields and free parameters, and implying the existence of new, yet undiscovered particles, have been an active field of study for decades. These models contain several fields, yielding a spectrum of scalar bosons with different properties. Many of the extended models are constrained, but not excluded, by the measured properties of the 125 GeV boson.

In singlet extensions, an additional singlet field is postulated. In *real singlet models*, the Higgs sector is extended with an additional real scalar SU(2) singlet. The Higgs doublet and the additional singlet provide two massive scalars: one of them is a Higgs boson similar to the one in the SM, while the other can be stable, providing a particle suitable as a dark matter candidate, which is missing in the SM [51]. In *complex singlet models*, where the additional scalar field is complex, also a third scalar particle appears [52]. Also models with several singlets have been studied [53]. While these models are interesting from the dark matter point of view, they fail to introduce explicit or spontaneous CP violation, which could explain the observed matter-antimatter asymmetry. As the singlet models also do not predict the existence of any charged Higgs bosons, which are the focus of this thesis, we do not discuss them further here.

Another simple way to extend the SM Higgs sector is to add one or several $\mathrm{SU(2)}_L$ doublets in the Higgs sector. The minimal models, where one new doublet is added, are called *two Higgs-doublet models* (2HDMs). The 2HDMs were first introduced in Ref. [17] in 1973, and since then they have been subject of great interest for several reasons. Firstly, the 2HDMs enter as a low-energy limit from several BSM theories, both supersymmetric and non-supersymmetric. Therefore they are useful as benchmark models used in designing experiments and interpreting their results, while the true nature of the wider theory remains unknown. Secondly, despite their simple structure, the 2HDMs predict a rich phenomenology with five scalar bosons, two of which carry electric charge, and the observation of these new particles could be within the reach of the LHC or its successors. Thirdly, the 2HDMs have potential to explain several open questions in physics. For example, some 2HDMs can yield a strong first-order phase transition, which would affect the electroweak baryogenesis and possibly explain the observed matter-antimatter asymmetry [17, 54, 55]. Some versions of the 2HDM, such as the so-called inert doublet model, can also provide a dark matter candidate [56]. Finally, the 2HDMs can satisfy certain theoretical constraints in a natural way, without excessive fine-tuning. For example, in the SM $\rho = m_{\mathrm{W}^\pm}^2 / (m_{\mathrm{Z}}^2 \cos^2 \theta_W)$ is equal to 1 at tree level (Eq. (2.22)). Experimentally, ρ has been confirmed to be very close to 1, and the 2HDMs automatically satisfy this relation at tree level [57].

As the 2HDMs are the simplest models that predict the existence of charged Higgs bosons, they serve as the main theoretical context for the work presented in this thesis. The properties of 2HDMs are discussed in detail in Section 3.1.1, and they are applied to predict the properties of charged Higgs bosons in Section 3.2. Also more complex models with multiple added doublets, and with higher representations such as triplets, have been developed. Unlike singlet and doublet models, the triplet models do not predict $\rho = 1$ at tree level. As the triplet models also predict the existence of charged Higgs bosons, their implications for the $H^\pm$ phenomenology are also briefly discussed in Section 3.2.

3.1.1 Two Higgs-doublet models

The 2HDM Higgs sector contains two complex $SU(2)_L$ doublets with hypercharge $Y = +1/2$, labeled Φ_1 and Φ_2. The most general renormalizable and gauge invariant Higgs potential that can be constructed with these two doublets is [18]

$$
\begin{aligned}
V(\Phi_1, \Phi_2) = & - \mu_1^2 \Phi_1^\dagger \Phi_1 - \mu_2^2 \Phi_2^\dagger \Phi_2 - (\mu_3^2 \Phi_1^\dagger \Phi_2 + \text{h.c.}) \\
& + \frac{1}{2}\lambda_1(\Phi_1^\dagger \Phi_1)^2 + \frac{1}{2}\lambda_2(\Phi_2^\dagger \Phi_2)^2 \\
& + \lambda_3(\Phi_1^\dagger \Phi_1)(\Phi_2^\dagger \Phi_2) + \lambda_4(\Phi_1^\dagger \Phi_2)(\Phi_2^\dagger \Phi_1) \\
& + (\frac{1}{2}\lambda_5(\Phi_1^\dagger \Phi_2)^2 + \lambda_6(\Phi_1^\dagger \Phi_1)(\Phi_1^\dagger \Phi_2) + \lambda_7(\Phi_2^\dagger \Phi_2)(\Phi_1^\dagger \Phi_2) + \text{h.c.}).
\end{aligned}
\tag{3.2}
$$

To suppress flavor-changing neutral currents (FCNCs) mediated by Higgs bosons, typically a Z_2 symmetry is postulated. Imposing this symmetry makes all fermions with a given charge to couple only to one doublet, and a theorem by Glashow and Weinberg states that in this case FCNCs are forbidden at tree level [58]. This prevents the FCNCs in a natural way, without additional fine-tuning.

In the following, we choose the Φ_1 (Φ_2) doublet as the even (odd) one under this symmetry, so they transform as $\Phi_1 \rightarrow \Phi_1$, $\Phi_2 \rightarrow -\Phi_2$. Imposing this symmetry forces $\lambda_6 = \lambda_7 = 0$. The remaining eight free parameters $\mu_1...\mu_3$ and $\lambda_1...\lambda_5$ can all be chosen to be real, ensuring CP-conservation in the Higgs sector. If μ_3 or λ_5 are allowed to be complex, CP-violation can occur. In the following we focus on the former case. The CP-violating versions of 2HDM are discussed further in Ref. [59]. In analogy to the single-doublet case discussed in Section 2.2, the potential $V(\Phi_1, \Phi_2)$ can be minimized

by non-zero vacuum expectation values

$$\langle\Phi_1\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v_1 \end{pmatrix}, \qquad \langle\Phi_2\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v_2 \end{pmatrix}, \tag{3.3}$$

where v_1 and v_2 satisfy $v_1^2 + v_2^2 = v^2 \approx (246\,\text{GeV})^2$. As previously, we expand the potential around the minimum as

$$\Phi_1 = \begin{pmatrix} \phi_1^+ \\ \frac{1}{\sqrt{2}}(v_1 + \phi_1^0) \end{pmatrix}, \qquad \Phi_2 = \begin{pmatrix} \phi_2^+ \\ \frac{1}{\sqrt{2}}(v_2 + \phi_2^0) \end{pmatrix}, \tag{3.4}$$

breaking the $\text{SU}(2)_L \times \text{U}(1)_Y$ symmetry into a residual $\text{U}(1)_{\text{EM}}$ symmetry. Working out the mass eigenstates, one finds out that the 2HDMs predict five different Higgs bosons: two neutral CP-even particles h and H, one neutral CP-odd particle A, and two charged Higgs bosons $\text{H}^\pm$ [18]. These particles arise as mixtures of the two doublets:

$$\begin{aligned}
h &= -\sin\alpha\,\text{Re}(\phi_1^0) + \cos\alpha\,\text{Re}(\phi_2^0), \\
H &= \cos\alpha\,\text{Re}(\phi_1^0) + \sin\alpha\,\text{Re}(\phi_2^0), \\
A &= -\text{Im}(\phi_1^0)\sin\beta + \text{Im}(\phi_2^0)\cos\beta, \\
H^+ &= -\phi_1^+\sin\beta + \phi_2^+\cos\beta,
\end{aligned} \tag{3.5}$$

where the angle α is the mixing angle of the CP-even states h and H, and the angle β is defined by the ratio of the vacuum expectation values:

$$\tan\beta = \frac{v_2}{v_1}. \tag{3.6}$$

At tree level, the masses of the corresponding particles are:

$$\begin{aligned}
m_{\text{H,h}}^2 &= \frac{1}{2}\left(\mathcal{M}_{11} + \mathcal{M}_{22} \pm \sqrt{(\mathcal{M}_{11} - \mathcal{M}_{22}^2 + 4\mathcal{M}_{12}^2)}\right), \\
m_{\text{A}}^2 &= \frac{2\mu_3}{\sin 2\beta} - \lambda_5 v^2, \\
m_{\text{H}^\pm}^2 &= m_{\text{A}}^2 + \frac{1}{2}v^2(\lambda_5 - \lambda_4),
\end{aligned} \tag{3.7}$$

where $\mathcal{M}$ is the CP-even mass-squared matrix diagonalized by a rotation by the angle α. By convention, $m_{\text{h}} \leq m_{\text{H}}$, so h picks the minus sign in the first equation. The eight parameters in $\mu_1...\mu_3$ and $\lambda_1...\lambda_5$, together with the constraint $v_1 + v_2 = v^2$, leave us

with seven free parameters. In the context of charged Higgs boson studies, these are often chosen as the four Higgs boson masses m_h, m_H, $m_{H^\pm}$ and m_A, the two angles α and β, and the scale parameter μ_3 that controls the mixing of the two doublets.

Obviously, the division between CP-even and CP-odd neutral particles arises only in the CP-conserving version of the theory. In the CP-violating versions of the theory three neutral states and two charged mass eigenstates are still present, but the three neutral states do not have definite CP quantum numbers. Even though the CP-conserving version is typically used as the benchmark, the $H^\pm$ searches are also sensitive to the CP-violating versions of 2HDM.

The Z_2-symmetric 2HDMs are classified into four different types [60], depending on how quarks and leptons couple to the two Higgs doublets. In Type I, Φ_2 couples to all fermions but Φ_1 does not. In Type II, Φ_1 couples to down-type quarks and charged leptons, while Φ_2 couples to up-type quarks. In Type X ("lepton-specific") 2HDMs, Φ_1 couples to charged leptons, whereas Φ_2 couples to all quarks. Finally, in Type Y ("flipped") 2HDMs, Φ_1 couples to down-type quarks, while Φ_2 couples to up-type quarks and charged leptons. These four types are summarized in Table 3.1.

In 2HDMs, the Higgs bosons couple to fermions via fermion-mass dependent Yukawa couplings similar to Eq. (2.23). The couplings between the $H^\pm$ and the fermions are given by

$$
\begin{aligned}
g_{H^+\bar{u}_i d_j} &= \frac{1}{\sqrt{2}} \frac{V_{ij}^{\text{CKM}}}{m_{W^\pm}} (\lambda_u \frac{1-\gamma_5}{2} + \lambda_d \frac{1+\gamma_5}{2}), \\
g_{H^+\bar{\nu}_\ell \ell} &= \frac{1}{\sqrt{2}} \frac{1}{m_{W^\pm}} \lambda_\ell \frac{1+\gamma_5}{2},
\end{aligned}
\tag{3.8}
$$

where the index i (j) runs over all up-type (down-type) quark flavors, $\ell = e, \mu, \tau$. The chirality operator γ_5 acts such that $(1-\gamma_5)/2$ $((1+\gamma_5)/2)$ selects the left-handed (right-handed) chiral states. The values of λ_u, λ_d and λ_ℓ depend on the 2HDM type as shown in Table 3.1. Hence for example the decay modes depend on the 2HDM type, as will be elaborated in Section 3.2.

Regarding gauge boson couplings, 2HDMs introduce a set of completely new tree-level couplings compared to the SM, of type $H^\mp W^\pm H_i$ and VVH_i, where H_i can be h, H or A and V refers to $W^\pm$ or Z. On the other hand the $H^\pm W^\mp Z$ coupling vanishes at tree level in 2HDMs. In fact, this is a general feature of models containing only Higgs singlets and doublets [57].

Table 3.1: The classification of Z_2-symmetric 2HDMs.

Type	u	d	ℓ	λ_u	λ_d	λ_ℓ
I	Φ_2	Φ_2	Φ_2	$\cot\beta$	$\cot\beta$	$\cot\beta$
II	Φ_2	Φ_1	Φ_1	$\cot\beta$	$-\tan\beta$	$-\tan\beta$
X	Φ_2	Φ_1	Φ_2	$\cot\beta$	$-\tan\beta$	$\cot\beta$
Y	Φ_2	Φ_2	Φ_1	$\cot\beta$	$\cot\beta$	$-\tan\beta$

In the *alignment limit* $\cos(\beta - \alpha) = 0$ ($\sin(\beta - \alpha) = 1$), the couplings of the lighter CP-even scalar h to fermions and gauge bosons agree with the SM prediction. Similarly, when $\sin(\beta - \alpha) = 0$ ($\cos(\beta - \alpha) = 1$), the heavier CP-even scalar H has SM-like couplings. Hence either h or H can correspond to the observed 125 GeV boson, the couplings of which have been confirmed to agree with the SM prediction. From this point of view, another interesting limit is the *decoupling limit*, where additionally $m_A \gg v^2$. In this limit, both H and $H^\pm$ become almost degenerate in mass with A, so all these masses are well above the electroweak scale, and do not largely affect the electroweak-scale physics. This way, the Higgs sector as provided in the SM arises naturally as a low-energy limit of the 2HDM.

In the very start of our discussion, the Z_2 symmetry was imposed to prevent tree-level FCNCs. Also versions of 2HDMs have been suggested where the tree-level FCNCs enter, but they are suppressed by a separate mechanism. For example, in so-called BGL models, the neutral currents are proportional to CKM matrix elements, so the small off-diagonal elements can suppress FCNCs to a tolerable level [61].

3.1.2 Supersymmetry

Supersymmetric models have for a long time constituted a popular class of BSM models. In supersymmetric models, each SM fermion (boson) is associated with a superpartner (an "sparticle"), which is a boson (fermion). If supersymmetry were an exact symmetry, the SM particles and their superpartners are identical in mass. However, if supersymmetry is broken (for example, broken spontaneously in analogy to the spontaneous breaking of the electroweak symmetry), the sparticles can be heavier than the SM particles and thus yet undiscovered.

While to date there is no experimental evidence to support supersymmetry, the theory has several appealing properties. Firstly, it solves the hierarchy problem in a natural

way, as the quadratic divergences arising from the terms containing SM particles are canceled by similar terms corresponding to their superpartners. Secondly, if an additional symmetry known as R-parity is included, preventing the sparticles from decaying into SM particles, the lightest supersymmetric particle constitutes a dark matter candidate. Thirdly, supersymmetry modifies the running of the weak, electromagnetic and strong couplings so that they unify into one in high energy scales. Additionally, many candidate theories that aim to integrate gravity with the other interactions, such as supergravity and superstring theories, are supersymmetric. Here we concentrate on the Higgs sector of the supersymmetric theories. A more comprehensive discussion of the supersymmetric models is given e.g. in Ref. [62].

In supersymmetric theories, each scalar field is associated with fermions of a given chirality, so two $SU(2)_L$ doublets are needed in the Higgs sector to provide masses to all fermions. In addition, a singlet scalar field can be added. In the Minimal Supersymmetric Standard Model (MSSM) the singlet field is not included, so the MSSM Higgs sector is a special case of Type II 2HDM [34].

Supersymmetry imposes additional constraints compared to more general 2HDMs, so only two parameters are needed to describe the MSSM Higgs sector at tree level (compared to seven parameters in our previous discussion). For $H^\pm$ studies, it is convenient to choose $m_{H^\pm}$ and $\tan\beta$ as the free parameters. The other Higgs boson masses are then determined by $m_{W^\pm}$ and m_Z as

$$
\begin{aligned}
m_A^2 &= m_{H^\pm}^2 - m_{W^\pm}^2, \\
m_{h,H}^2 &= \frac{1}{2}m_A^2 + m_Z^2 \pm \sqrt{(m_A^2 + m_Z^2)^2 - 4m_Z^2 m_A^2 \cos^2 2\beta}.
\end{aligned}
\tag{3.9}
$$

The mixing angle α is determined by

$$
\begin{aligned}
\cos 2\alpha &= -\frac{m_A^2 - m_Z^2}{m_H^2 - m_h^2}\cos 2\beta, \\
\sin 2\alpha &= -\frac{m_H^2 + m_h^2}{m_H^2 - m_h^2}\sin 2\beta,
\end{aligned}
\tag{3.10}
$$

and the mixing parameter μ_3 is fixed to

$$
\mu_3^2 = m_A^2 \frac{\tan\beta}{1 + \tan^2\beta}.
\tag{3.11}
$$

Eq. (3.9) implies that, at tree level, $m_\text{h} < m_\text{Z}$, while $m_\text{H} \geq m_\text{Z}$, $m_\text{A} \geq m_\text{h}$ and $m_{\text{H}\pm} \geq m_{\text{W}\pm}$. These relations are, however, largely modified by *radiative corrections* from higher orders of perturbation. Especially the loop processes that involve top quarks alter the mass relations, so that it is completely feasible to obtain a mass of 125 GeV for the lighter CP-even neutral boson h.

As the radiative corrections depend on the properties of all particles present in the theory, they introduce dependence on several parameters in the MSSM beyond $m_{\text{H}\pm}$ and $\tan\beta$. To study the phenomenological consequences of different parameters, a number of *benchmark scenarios* have been proposed. The most important parameters for the Higgs sector corrections are known as X_t and M_SUSY. In this thesis, the $\text{H}^\pm$ search results are interpreted in the so-called MSSM $m_\text{h}^{\text{mod-}}$ scenario. This scenario is specified using low-energy MSSM parameters and it is designed to give a mass of approximately 125 GeV for the light CP-even Higgs boson over a wide region of the parameter space. The parameter X_t (M_SUSY) is given a large negative (positive) value to obtain a satisfactory agreement with flavor physics observables. Another variant of this scenario, called $m_\text{h}^{\text{mod+}}$, is defined identically except that the parameter X_t is given a large positive value, yielding a better agreement with the muon magnetic moment measurements. The exact definitions and detailed discussions of these and other modern benchmark scenarios can be found in Ref. [63, 64].

3.2 Properties of charged Higgs bosons

Charged Higgs bosons are predicted in many different models containing at least two Higgs doublets, or higher representations such as triplets. As the (CP-conserving) 2HDMs are the minimal extension of the SM where $\text{H}^\pm$ appears, they serve as good benchmarks for experimental searches. Therefore the $\text{H}^\pm$ production processes and decays channels are discussed here in the context of 2HDMs, including the special case of the MSSM.

In multiple-doublet models with ≥ 3 Higgs doublets, each added doublet introduces another $\text{H}^\pm$ pair in the particle spectrum [65]. As discussed in Section 3.1.1, in doublet models the charged Higgs bosons do not couple to vector boson pairs at tree level. However, if the Higgs sector contains higher-order representations such as triplets, couplings such as $\text{H}^\pm$ W Z occur already at tree level and can alter the production and decay modes significantly compared to the doublet models [66–68]. While the triplet models and their phenomenology are beyond the scope of this thesis, we return to

them briefly in Section 3.2.3 while discussing experimental constraints from different types of experiments.

3.2.1 Production mechanisms

Examples of leading order (LO) diagrams describing the $H^\pm$ production in 2HDM in different mass regions are shown in Figure 3.1.

Light charged Higgs bosons, with a mass smaller than the mass difference between the top and bottom quarks ($m_{H^\pm} < m_t - m_b$), are predominantly produced in decays of top quarks ($t \to bH^\pm$), as shown in Figure 3.1 (left). As the top quarks are mostly produced in pairs at the LHC, this production mode is called *double-resonant* production.

The *heavy* charged Higgs bosons, with $m_{H^\pm} > m_t - m_b$, are dominantly produced in association with a top quark, as shown in Figure 3.1 (middle). We denote this production mode by $pp \to tbH^\pm$. At finite order, this production process can be understood in two complementary ways. In the *four-flavor scheme* (4FS), no b quarks are present in the initial state. The gluon fusion ($gg \to tbH^\pm$) is the dominant production mode, and the quark-antiquark annihilation ($qq' \to tbH^\pm$) the secondary one. Gluon splittings can create $b\bar{b}$ pairs which are nearly collinear to the initial gluon. In the *five-flavor scheme* (5FS), these gluon splittings are summed in all orders of perturbation by introducing bottom parton densities. This way the b quark can be present already in the initial state, and in fact $gb \to tH^\pm$ becomes the dominant process at LO. In 5FS, gluon fusion and qq' annihilation enter at NLO.

Figure 3.1: Leading order diagrams describing the charged Higgs boson production. Double-resonant top quark production (left) is the dominant process for light $H^\pm$, whereas the single-resonant top quark production (middle) dominates for heavy $H^\pm$. For the intermediate mass region ($m_{H^\pm} \sim m_t$), both production modes and their interplay with the nonresonant top quark production (right) must be taken into account. Charge-conjugated processes are implied.

Figure 3.2: $H^\pm$ production cross sections calculated in 4FS (red line) and 5FS (blue line) and the Santander-matched prediction (black line) for $m_{H^\pm} = 200\,\text{GeV}$ (left) and $2000\,\text{GeV}$ (right). The dashed lines (green band) show the theoretical uncertainties for the separate (matched) cross sections. [69]

The 4FS and 5FS can be shown to agree when all orders of perturbation are included in the calculations, but at finite order significant differences may arise. To understand these differences, calculations have been carried out in both schemes at LO and NLO, comparing both differential and inclusive $H^\pm$ cross sections [69]. While at LO large differences have been found, at NLO results are typically compatible within the uncertainties. The dominant sources of theoretical uncertainty, however, are different: the 4FS results are most affected by the chosen renormalization and factorization scales, while the 5FS results depend more on the chosen parton distribution functions and their uncertainties. To benefit from both approaches and reduce the overall uncertainty, an approach known as *Santander matching* has been adopted to extrapolate between the cross sections calculated with the two schemes [69]:

$$\sigma_{\text{matched}} = \frac{\sigma_{4\text{FS}} + w\sigma_{5\text{FS}}}{1 + w}, \tag{3.12}$$

where weight w depends logarithmically on $m_{H^\pm}$:

$$w = \ln \frac{m_{H^\pm}}{mb} - 2. \tag{3.13}$$

The Santander matching is consistent with asymptotic behavior in limits $m_{H^\pm}/m_b \to 1$ and $m_{H^\pm}/m_b \to \infty$ where the two schemes are known to agree. In Figure 3.2, examples of the $H^\pm$ cross sections are shown, calculated both in 4FS and 5FS and then matched using Eq. (3.12). In the calculation, $m_{H^\pm}$ is set to $200\,\text{GeV}$ (left) and $2\,\text{TeV}$

Figure 3.3: $H^{\pm}$ production cross sections for the intermediate mass range as obtained from a full NLO calculation of the $pp \to H^{\pm}W^{\mp}b\bar{b}$ process (solid lines), compared to the exclusive calculations of $t \to bH^{\pm}$ and $pp \to tbH^{\pm}$ processes (dashed lines). [70]

(right). In simulation of signal samples for experimental searches, the usage of 4FS is recommended since it is found to produce more accurate differential distributions [69].

In the intermediate region near the mass of the top quark ($m_{H^{\pm}} \sim m_t$), the nonresonant top quark production mode (Figure 3.1, right) also contributes and the full process $pp \to H^{\pm}W^{\mp}b\bar{b}$ must be calculated to correctly account for all three production mechanisms and their interference. For a long time, the lack of reliable cross section calculations for the full process prevented the interpretations of experimental results in the intermediate mass region. In 2017, the first rigorous NLO cross section calculation for the full process $pp \to H^{\pm}W^{\mp}b\bar{b}$ in the intermediate region was published [70].

Figure 3.4: Branching fractions for $H^{\pm}$ decaying into different final states as a function of the $H^{\pm}$ mass, calculated for 2HDM Types I and X (left), and for Types II and Y (right), assuming $\sin(\beta - \alpha) = 1$ and $\tan \beta = 30$. [71]

The results are shown in Figure 3.3. It is reassuring to notice that the calculated cross section (solid line) agrees with the standard results for $t \rightarrow bH^{\pm}$ ($pp \rightarrow tbH^{\pm}$) when $m_{H^{\pm}}$ is well below (above) the top quark mass. The kink in the cross section corresponds to the kinematic threshold $m_{H^{\pm}} = m_t - m_b$.

3.2.2 Decay modes

The $H^{\pm}$ *branching fractions*, which define the probability of each possible decay mode, depend on the mass of $H^{\pm}$, which limits the kinematically allowed decays and constrains the phase space of the decay products. The branching fractions also depend on the strength of the coupling between $H^{\pm}$ and the decay products. As the Yukawa couplings of Higgs bosons to fermions increase linearly with the fermion mass, the decays to third-generation quarks and leptons dominate the $H^{\pm}$ decays over the first and second generation. Since the couplings depend on the 2HDM type, the branching fractions are model-dependent.

Figure 3.4 shows an example of branching fractions as a function of $m_{H^{\pm}}$, as predicted in different types of 2HDM in the alignment limit $\sin(\beta - \alpha) = 1$ and assuming $\tan \beta = 30$ [71]. For the light $H^{\pm}$, we notice that in Types II and X the $H^{\pm}$ decays almost exclusively to a tau lepton and a neutrino ($H^{\pm} \rightarrow \tau^{\pm} \nu_{\tau}$) while in Types I and

Figure 3.5: Branching fractions for $H^\pm$ decaying into different final states as a function of the $H^\pm$ mass, calculated in the MSSM $m_h^{\text{mod-}}$ benchmark scenario assuming $\tan\beta = 10$ (left) and $\tan\beta = 50$ (right). [64]

Y also decays into cs or cb can be important. For the heavy $H^\pm$, the decays into tb or into $HW^\pm$ dominate, depending on the model.

In the MSSM, the 2HDM Type II calculations need to be modified to incorporate the effects from supersymmetric particles. In Figure 3.5, the MSSM branching fractions are shown as a function of $m_{H^\pm}$ assuming the $m_h^{\text{mod-}}$ benchmark scenario [64]. The dependence on $\tan\beta$ is illustrated by showing the results for $\tan\beta = 10$ (left) and $\tan\beta = 50$ (right). Again, the branching fraction to the $\tau^\pm \nu_\tau$ final state is almost 100% for the light $H^\pm$. Above the top quark mass, the tb decay channel opens up, but the branching fraction to $\tau^\pm \nu_\tau$ remains sizable.

In this thesis, we focus on the $H^\pm \to \tau^\pm \nu_\tau$ decay channel. While in the case of light $H^\pm$ this is obviously an interesting channel, it is less evident for the heavy mass region by looking at Figures 3.4 and 3.5. We should note, however, that in 2HDM Types II and Y, the coupling of the $H^\pm$ to leptons is proportional to $\tan\beta$, so the branching fraction to a tau lepton and a neutrino can be significant for large values of $\tan\beta$. Secondly, the $\tau^\pm \nu_\tau$ final state is experimentally attractive compared to e.g. the tb final state, since it offers several effective handles to discriminate the $H^\pm$ signal from the SM background. These experimental aspects will be discussed later in the third part of this thesis.

3.2.3 Experimental constraints

The evolution of the BSM models that predict the existence of charged Higgs bosons is partially guided by experimental results that can be used to falsify a model or, more commonly, some part of its parameter space. Here we review the latest experimental constraints on the BSM Higgs sector, focusing on the 2HDMs and the charged Higgs bosons predicted by them.

The direct searches for charged Higgs bosons at colliders, including the topic of this thesis, have the benefit of providing robust and essentially model-independent results. The non-observation of the $H^{\pm}$ at the LHC and in the previous colliders such as LEP and Tevatron limits the mass of the $H^{\pm}$ and constrains the models that predict them [19–21]. Prior to the LHC, the most stringent limits have been set by the LEP experiments. At LEP, $H^{\pm}$ were searched for assuming pair production $e^{+}e^{-} \rightarrow \gamma/Z \rightarrow H^{+}H^{-}$. The results exclude $H^{\pm}$ with mass below 80 GeV when the $H^{\pm} \rightarrow W^{\pm}h$ decay is not present, or up to 72.5 GeV when this decay is included and $m_{h} > 12$ GeV [19].

During the LHC Run 1, the search efforts focused mostly on light $H^{\pm}$. The searches by the CMS and ATLAS Collaborations covered the $\tau^{\pm}\nu_{\tau}$ [22–24, 26, 27], cs [25, 29], and cb [33] decay modes using the data collected at $\sqrt{s} = 7$–8 TeV. The $\tau^{\pm}\nu_{\tau}$ results, however covered also the heavy $H^{\pm}$ scenario, complemented by heavy-$H^{\pm}$ specific searches in the tb channel [27, 28].

Since the start of the Run 2 in 2015, the increased center-of-mass energy of $\sqrt{s} = 13$ TeV and the higher instantaneous luminosities, implying higher amounts of collected data, have allowed kinematic access to heavy-$H^{\pm}$ scenarios, and the focus has shifted to this region. This thesis focuses on the search in the $\tau^{\pm}\nu_{\tau}$ final state at $\sqrt{s} = 13$ TeV, based on the data collected by the CMS experiment in 2016 [72, 73]. The ATLAS Collaboration has performed similar searches based on data collected in 2015–2016 [30, 31]. In addition, both experiments have covered the tb decay channel at 13 TeV [32, 74, 75]. These recent searches have probed possible $m_{H^{\pm}}$ hypotheses up to several TeV.

While all the above searches assume $H^{\pm}$ production either via $t \rightarrow bH^{\pm}$ or $pp \rightarrow tbH^{\pm}$, as motivated by the 2HDMs, also production modes beyond the 2HDM have been included in the searches. For example, in triplet models such as Georgi-Machacek models, the $H^{\pm}W^{\pm}Z$ coupling is allowed at tree level, meaning that the $H^{\pm}$ can be produced in vector boson fusion and that it decays dominantly into $W^{\pm}Z$ [67]. Search motivated by this scenario have been performed by the ATLAS and CMS
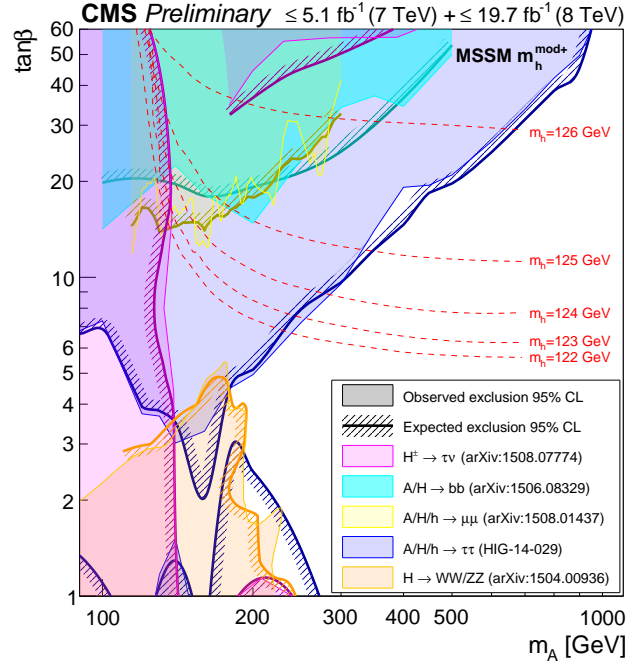
Figure 3.6: Exclusion limits from different BSM Higgs boson searches performed by the CMS Collaboration using data from the LHC Run 1, interpreted in the MSSM $m_h^{\text{mod+}}$ benchmark scenario. The white area represents the parameter space not excluded by the results. The area below the dashed red line with text $m_h = 125\,\text{GeV}$ is excluded provided that the observed 125 GeV Higgs boson corresponds to the lightest CP-odd scalar h of the MSSM. [81]

Collaborations [76–78]. In triplet models, also doubly-charged Higgs bosons appear, and they have been searched for by the ATLAS and CMS Collaborations [79, 80].

In addition to charged Higgs boson searches, a wide search program is ongoing at the LHC to discover additional neutral Higgs bosons, such as the CP-even and CP-odd ones predicted in the 2HDMs. Some of these searches are motivated by general 2HDMs, while others target specifically the MSSM. The exclusion limits obtained from these searches can be used to constrain the $H^\pm$ mass in any model that predicts a relation between $m_{H^\pm}$ and a neutral scalar boson mass. For example, in the MSSM the masses of the h, H and A bosons are connected to $m_{H^\pm}$ via Eq. (3.9), allowing reinterpretations of the neutral boson search results in $(m_{H^\pm}, \tan\beta)$ plane.

The ATLAS and CMS Collaborations have performed a large number of searches for neutral BSM Higgs bosons, in decay channels where a neutral BSM boson decays into a final state of two SM particles, such as $\tau\tau$ [82,83], $b\bar{b}$ [84], $t\bar{t}$ [85], $\mu\mu$ [86,87], $\gamma\gamma$ [88,89], $W^\pm W^\mp$ [90,91], and ZZ [92,93]. Additionally, decay channels involving multiple

Higgs bosons, such as $A \rightarrow Zh$ have been studied in several final states [94–96], as well as a scenario where the observed 125 GeV boson decays into two lighter ones ($H \rightarrow hh$) [97–99]. A summary of constraints in the MSSM $m_h^{\text{mod}+}$ benchmark scenario arising from the CMS search results for neutral and charged Higgs bosons is shown as a function of m_A and $\tan\beta$ in Figure 3.6. The $A/H/h \rightarrow \tau\tau$ is found to exclude almost half of the plane, complemented by limits from other channels.

The limits for neutral Higgs bosons can have exclusion power in the regions of parameter space that are not kinematically accessible by direct $H^{\pm}$ searches, or where the direct searches are not very sensitive for example due to difficult backgrounds. The interpretations of neutral Higgs boson results are, however, always strongly model-dependent. This is illustrated in Figures 3.7 and 3.8, which show exclusion limits as interpreted in the ($m_{H^{\pm}}$, $\tan\beta$) plane of Type I (left) and Type II (right) 2HDMs. In Figure 3.7, the masses are assumed to depend on each other in an MSSM-inspired way, while in Figure 3.8, $m_H = m_A = 1$ TeV. In both cases, m_h is fixed to 125 GeV. The green area indicates the parameter range allowed by the LEP and LHC searches, and the colored regions correspond to parameter values excluded by results from collider experiments. Comparing the two figures, we observe that in the first case the neutral boson searches in the $\tau\tau$ (grey) and $\gamma\gamma$ (orange) channels complement the direct searches, such as the $H^{\pm} \rightarrow \tau^{\pm}\nu_{\tau}$ search channel (red). However, in the second case where the H and A are heavy, the searches for them hardly impose any constraint on the $H^{\pm}$ mass, while the direct $H^{\pm}$ searches exclude the same regions as in the first case. Thus as long as we do not know which type of 2HDM is correct, if any, direct searches for $H^{\pm}$ remain important.

The properties of the observed Higgs boson set constraints to the extended Higgs sector. The mass of the observed boson has been measured to be 125.09 ± 0.21 (stat) ± 0.11 (syst) GeV [40], and the signal rates associated with different production modes and final states have been found to agree with the SM expectations [42]. Therefore any model with an extended Higgs sector needs to provide such a boson, consistent with the measured couplings. For example the $h \rightarrow \gamma\gamma$ decay rate can be significantly altered by the presence of a light $H^{\pm}$ loop [65]. The models that predict relations between the masses of different Higgs bosons, such as MSSM are heavily constrained by the observed 125 GeV boson, whereas theories with larger number of free parameters, such as generic 2HDM models, can accommodate the observed boson more easily.

The combined results from the ATLAS and CMS Collaborations suggest that the couplings of the observed 125 GeV boson are close to the SM prediction [42]. As
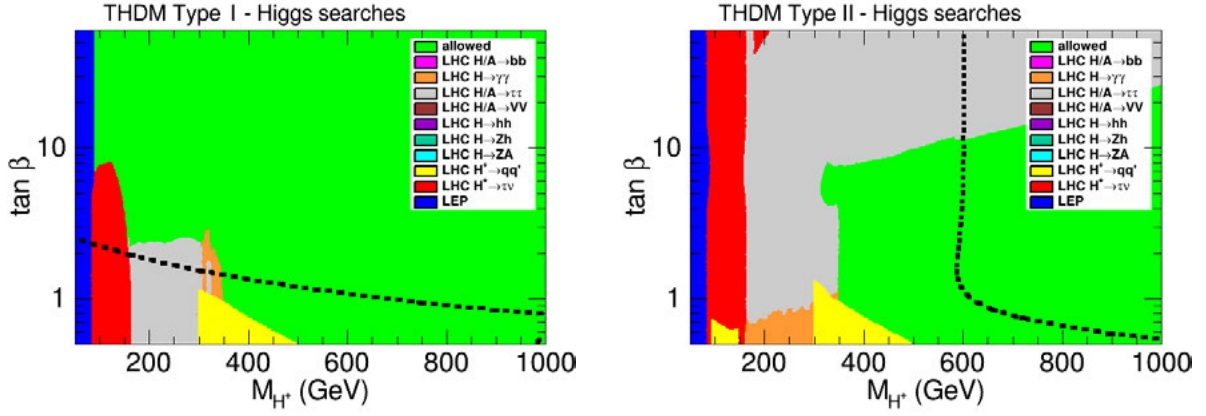
Figure 3.7: Constraints from LEP and LHC experiments interpreted in the $(m_{H^\pm}, \tan\beta)$ plane of Type I (left) and Type II (right) 2HDMs, in a scenario where $m_{H^\pm}$ is related to the neutral boson masses in an MSSM-like way. The regions below and left of the dashed line are excluded by flavor physics measurements. [71]

discussed in Section 3.1.1, in 2HDM such couplings for the h boson are achieved near the alignment limit where $\cos(\beta - \alpha) = 0$, without constraining $m_{H^\pm}$. Hence the measured couplings of the 125 GeV boson force $\cos(\beta - \alpha)$ to be small. The allowed range of variation around zero depends on the 2HDM type and on $\tan\beta$, as detailed in Ref. [100]. In Type I, which is the least constrained one, the observed couplings suggest that $|\cos(\beta - \alpha) < 0.4|$.

Results of electroweak precision measurements can also be used to constrain the 2HDMs. For example, the measured values of Peskin-Takeuchi parameters [101, 102] constrain the possible mass hierarchy of H, A and $H^\pm$ bosons, requiring that either $m_A \approx m_{H^\pm}$ or $m_H > m_{H^\pm}$ [100]. The existence of $H^\pm$ would also show up indirectly apparent violations of lepton universality in processes mediated by $W^\pm$ [103, 104].

The measurements of the anomalous moment of muon, $a_\mu = (g_\mu - 2)/2$ have tension with respect to the SM prediction [35]. As a_μ is sensitive to the coupling between muon and photon, its value could be modified by additional scalar bosons in the loops. Therefore the observed values can introduce constraints in 2HDMs, excluding $\tan\beta < 3$ (9) for $m_{H^\pm} = 200$ GeV (1 TeV) in Types II and X [100]. In other 2HDM types the constraints are weaker.

In addition to collider searches, the flavor physics measurements at the LHC and in B meson factories such as Belle provide important constraints on $H^\pm$ properties, since the $H^\pm$ is expected to enter all tree-level and loop-level decay processes of the

Figure 3.8: Constraints from LEP and LHC experiments interpreted in the $(m_{H^\pm}, \tan\beta)$ plane of Type I (left) and Type II (right) 2HDMs, in a scenario where $m_{H^\pm}$ is independent of neutral boson masses and $m_H = m_A = 1\,\text{TeV}$. The regions below and left of the dashed line are excluded by flavor physics measurements. [71]

B mesons where $W^\pm$ is present. The presence of $H^\pm$ can alter the results significantly, leading to observable deviations from the SM predictions.

In 2HDMs, the effect on the flavor physics observables mostly depends on $m_{H^\pm}$, the 2HDM type and $\tan\beta$ (all of which affect the Yukawa couplings), while the dependence on the other parameters in the model is small. Hence 2HDMs provide a convenient framework to interpret also the flavor-physics constraints on possible BSM Higgs sectors. Among several processes and observables modified by the presence of additional Higgs bosons, the strongest constraints are provided by measurements of the branching fraction $\mathcal{B}(B \to X_s \gamma)$. In Type II 2HDMs, also $\mathcal{B}(B \to \tau^\pm \nu_\tau)$ contributes significantly. In this process, $H^\pm$ appears at tree level, interfering with $W^\pm$ destructively and hence reducing the branching fraction compared to the SM prediction. Other sources of constraints include $B_s \to \mu\mu$, $B_s \to D^{(*)}\tau^\pm\nu_\tau$, and Δm_{Bs} [71, 100].

The combined effect of constraints from several flavor physics processes is marked in Figures 3.7 and 3.8 as thick dashed lines. In Type I 2HDMs, the flavor constraints exclude only very small values of $\tan\beta$ across the $m_{H^\pm}$ mass range. In Type II, the constraints are much stronger, excluding $H^\pm$ masses below $600\,\text{GeV}$ for all $\tan\beta > 1$. In Type X (Y), the combined effect of flavor physics constraints is similar to that in Type II (I). In addition, recent measurements on $B_s \to \mu\mu$ exclude a region in the top-right corner of the Type II plot, above $\tan\beta$ values of 15–25 for $m_{H^\pm}$ values from $600\,\text{GeV}$ to $1\,\text{TeV}$ [100].

To summarize, constraints on the $H^\pm$ properties, typically interpreted in the context of Z_2-symmetric CP-conserving 2HDMs, arise from collider searches for charged and neutral Higgs bosons, from the measured properties of the 125 GeV Higgs boson, from electroweak precision measurements and from flavor physics experiments. The model-independent limits are still relatively low, and charged Higgs bosons as light as ≈ 75 GeV are still compatible with the available experimental data. In specific scenarios, significantly more stringent constraints arise. For example, in Type II 2HDMs the strongest constraints are given by the flavor observables which suggest that $m_{H^\pm}$ must be larger than 600 GeV. While interpreting the constraints from different types of experiments in simple benchmark scenarios, such as different types of 2HDMs, is useful for understanding the interplay of different parameters of the SM and of the chosen BSM model, direct collider searches for $H^\pm$ are the only way to obtain robust, essentially model-independent results.

PART II

Experimental methods

Chapter 4

Large Hadron Collider

4.1 Overview

The Large Hadron Collider (LHC) is the largest and most powerful particle accelerator ever built. The project was proposed already in 1990, approved in 1994, and the collider started operating in 2008. The LHC is situated at CERN, Geneva, across the border between Switzerland and France. In this chapter, the main properties of the LHC and some key concepts of proton-proton (pp) collisions are discussed. A detailed description of the LHC machine can be found in Ref. [13].

The LHC is placed in a circular tunnel with a circumference of 26 659 m. The tunnel was previously occupied by the Large Electron-Positron Collider (LEP), a machine that operated from 1989 to 2000, reaching a maximum center-of-mass energy of 209 GeV. Large particle colliders are placed in tunnels because the bedrock provides good shielding against external radiation that might bias the measurements, while also absorbing the ionizing radiation produced by the collider. The LHC tunnel is situated between the Jura mountains and Lake Geneva, 50–175 meters underground and inclined towards Lake Geneva by 1.4%. Its location and the tilt are determined by geological considerations, combined with the aim to minimize the depth of the shafts used to access the tunnel, and the need to connect the LEP/LHC tunnel to the SPS pre-accelerator.

At the LHC, two proton beams circulate in opposite directions and collide in different *interaction points*, surrounded by detectors that record the collisions. The LHC is also used to collide lead ions with each other or with protons, but in the following discussion the focus will be on the pp collisions. The LHC accelerates protons instead of electrons (or positrons) because it is designed to reach as high collision energy as

possible. In a circular orbit charged particles lose energy as they emit *synchrotron radiation*, and this energy loss needs to be compensated by further acceleration to keep the collision energy constant, which consumes a lot of power. But since the synchrotron radiation is inversely proportional to the fourth power of the particle mass, the synchrotron radiation from protons is negligible compared to electrons.

The energy available to produce new particles in the collisions is the *center-of-mass energy*, denoted as $\sqrt{s}$, where $s = (p_1 + p_2)^2$ is a Mandelstam variable calculated from the four-momenta p_1 and p_2 of the two colliding protons. Thus in the collisions of two proton beams, the center-of-mass energy is simply the sum of the beam energies ($2E_{\text{beam}}$). On the other hand, in fixed-target experiments, the center-of-mass energy is proportional to $\sqrt{E_{\text{beam}}}$, which makes it more difficult to reach large center-of-mass energies.

The motivation for high collision energies can be seen from Figure 4.1, which shows a theoretical prediction [105] for the production cross sections for different processes as a function of the center-of-mass energy $\sqrt{s}$. For most processes of interest, such as the production of $W^{\pm}$ and Z bosons and Higgs bosons, the cross section and thus the production rate at the LHC increases with $\sqrt{s}$, so maximizing $\sqrt{s}$ maximizes the rate at which these particles are produced at the collider. The top curve in Figure 4.1 shows the total inelastic pp cross section (σ_{tot}). It is dominated by soft QCD interactions, characterized by low *momentum transfer*, meaning that the particles produced in these interactions have low momentum. The production of heavy short-lived particles such as top quarks, vector bosons or Higgs bosons requires a *hard collision*, where two partons interact with a large momentum transfer. By comparing the cross sections of these specific processes (e.g. σ_W, σ_Z, σ_t) we notice that these processes are very rare among all interactions that make up the σ_{tot}. Therefore clever *trigger* algorithms are needed to identify these hard interactions in-situ in the experiments, as discussed later in Chapter 6.

In circular colliders such as the LHC, the collision energy is limited by the need to keep the particles in (approximately) circular orbits at the center of the beam pipes. For a particle with momentum $|\vec{p}|$ and charge Q, the condition for circular orbit is

$$|\vec{p}| = QRB, \tag{4.1}$$

where R is the radius of the orbit and B is the strength of the magnetic field. Thus after maximizing R with a long tunnel, the magnet technology is the limiting factor

Figure 4.1: The predicted total inelastic pp interaction cross section (σ_{tot}) and production cross sections for several specific types of events as a function of the center-of-mass energy $\sqrt{s}$. The discontinuity at $\sqrt{s} \approx 4\,\text{TeV}$ corresponds to a switch from proton-antiproton cross sections to proton-proton cross sections. [105]

for $|\vec{p}|$ and thus for $\sqrt{s}$. The superconducting LHC magnets are discussed later in Section 4.2.1.

The LHC was designed to reach a maximum center-of-mass energy of $\sqrt{s} = 14\,\text{TeV}$. During its operation, the collision energy has been increased gradually. After low-energy tests in 2008–2009, the LHC Run 1 started first at $\sqrt{s} = 7\,\text{TeV}$ in 2010–2011, and then the energy was raised to 8 TeV in 2012. For Run 2 in 2015–2018, 13 TeV energy was used. In the future, the collision energy will presumably be raised to the design value of 14 TeV, and the possibility of further raising it to 15 TeV, the maximum theoretically achievable with the current magnets, is also considered [106].
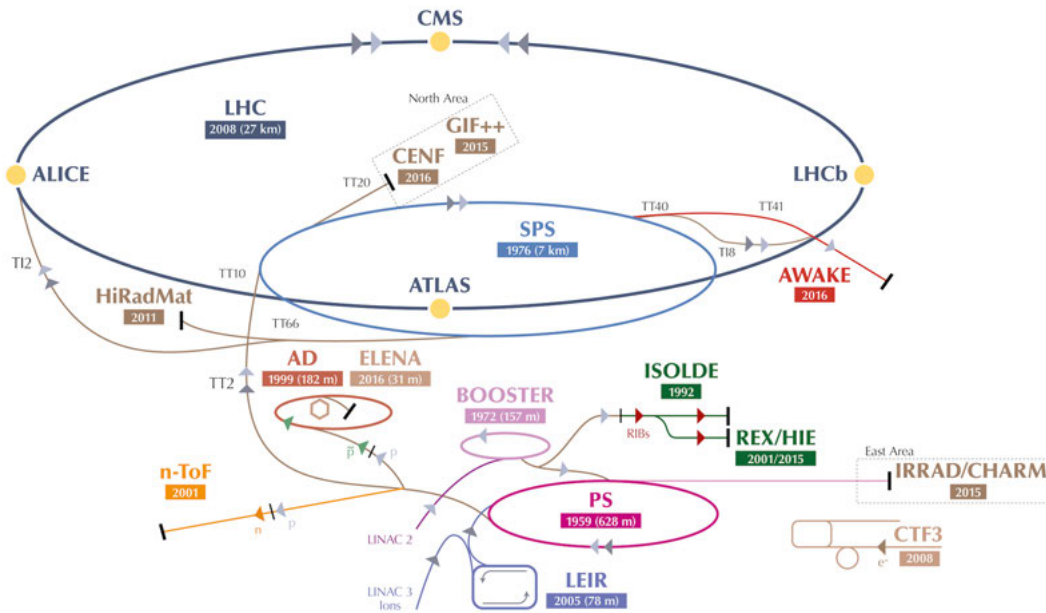
Figure 4.2: A sketch of the LHC and its pre-accelerators (LINAC2, Booster, Proton Synchrotron (PS) and Super Proton Synchrotron (SPS)), and different experimental facilities attached to them. For some of them, the year when the operation started and the tunnel length are also given. [107]

4.2 Architecture

A bottle of hydrogen is used as the proton source for the LHC. The hydrogen atoms are ionized using an electric field produced with cylindrical metal construction known as *Duoplasmatron*. The protons are first accelerated with a linear accelerator up to an energy of 50 MeV (corresponding to a velocity of $0.314c$), and then transferred into the Proton Synchrotron Booster (PSB), where they obtain an energy of 1.4 GeV ($0.916c$). After PSB, the beam is injected into the Proton Synchrotron (PS), which accelerates it to 25 GeV ($0.9993c$). Then the Super Proton Synchrotron (SPS) is then used to increase the energy up to 450 GeV ($0.999998c$). At this energy, the protons can be injected as two opposite-direction beams into the LHC, where it takes approximately 20 minutes to ramp up the energy from 450 GeV to the final collision energy, which is 6.5 TeV per beam during Run 2, and corresponds to $0.99999999c$. The LHC can also accelerate lead ions (Pb^{54+}), which can obtain an energy of 2.56 TeV per nucleon. An overview of the LHC, its pre-accelerators and the attached experimental facilities is shown in Figure 4.2.

In order to avoid collisions with gas molecules, the beams traverse the LHC in vacuum tubes. An ultrahigh vacuum of 10^{-13} atmospheric pressures is maintained inside the tubes. This corresponds roughly to the pressure of the "atmosphere" of the Moon.

If the fill is successful and the beam stays stable, it can circulate in the machine for more than ten hours, while the intensity of the beam gradually decreases due to the consumption of protons in the collisions at the interaction points and other smaller particle losses.

Each fill ends with a *beam dump*, when a set of dedicated magnets vertically deflect the beams to a separate 700 m long tunnel that ends with *beam dump blocks* made of graphite plates. Special *kicker magnets* are used to dilute the beams before they hit the graphite plates. To protect the accelerator and detector equipment, dedicated beam loss sensors continuously monitor the beam. If the beam unexpectedly becomes unstable, it is automatically dumped within 0.3 ms.

Even though the LHC is referred to as a "circular collider", the trajectory of the beam is not perfectly circular. It consists of eight straight *insertions* (each 545 m in length), combined by eight *arcs* (each 2.45 km in length). Each arc is occupied with 154 dipole magnets that bend the beam. Out of the eight insertions, four host the large physics experiments, one is occupied by the radio frequency cavities used for acceleration, two are reserved for beam cleaning and is one connected to the beam dumping tunnel. The whole structure needs to be precisely controlled for successful operation. For example, the gravity of the Moon needs to be taken into account in beam orbit corrections, since the LHC circumference changes by 1 mm as the Earth's crust rises during ground tide [108].

The four large experiments situated at the interaction points are ALICE, ATLAS, CMS and LHCb. ATLAS and CMS are *multipurpose experiments*, with wide physics programs covering different areas from Higgs boson measurements to b-hadron and heavy ion physics. ALICE is dedicated for studying heavy-ion collisions and the resulting quark-gluon plasma, while LHCb focuses on rare decays containing b and c quarks. Three smaller experiments are placed close to the large ones: LHCf (The Large Hadron Collider forward, using the LHC beam to understand cosmic rays) next to ATLAS, the Monopole and Exotics Detector MoEDAL next to LHCb, and TOTEM (TOTal cross section, Elastic scattering and diffraction dissociation Measurement at the LHC) next to CMS. These experiments, as well as other experimental facilities attached to the smaller accelerators, are also shown in Figure 4.2.

4.2.1 Magnets and beam control

As discussed above, the strength of the bending magnetic field sets the limit of the collision energy. The 1232 dipole magnets in the LHC arcs provide a magnetic field strong enough to keep the 6.5 TeV protons in the orbit. They are superconducting electromagnets, each 15 m long and 35 tonnes in weight, and they operate at a current of 11 kA, needed to create a magnetic field of 8 T. The coils are made of niobium-titanium (NbTi) cable, consisting of 20–30 thousand 7 μm thin filaments. In total, the LHC contains 7600 km of NbTi cable.

The LHC is divided into eight *sections*, each of which is powered independently. The dipole magnets of each sector are connected in series inside a shared cryostat. To achieve superconductivity, the magnets are cooled down to 1.9 K temperature with superfluid helium. Its high thermal conductivity makes it suitable for refrigeration of large superconducting systems. Cooling of the LHC magnets requires some 120 tonnes of helium.

As the protons are positively-charged, they repel each other electrically in the LHC beams. In total, 392 *quadrupole magnets* are used to constrain either the height or the width of the beam to keep it collimated. At the LHC, the magnets are placed in 23 *FODO cells*, each of which contains one horizontally focusing and one vertically focusing quadrupole magnet, with dipole magnets and higher-multipole magnets placed between them. The higher-multipole magnets, such as sextupole, octupole and decapole magnets, provide additional control and focusing of the beam. The overall effect of the magnet system is that the protons keep oscillating around the center of the vacuum tube. Near the interaction points surrounded by the experiments, *inner triplet magnets* are used to squeeze the beam from a transverse size of $\approx$200 μm down to $\approx$20 μm. When all types of magnets are taken into account, the LHC beam control system contains approximately ten thousand magnets.

4.2.2 Radio frequency cavities and beam acceleration

The protons are accelerated by 16 *radio frequency cavities* that provide a strong electric field of 5 MV/m. After reaching the maximum energy, they are designed to maintain the longitudinal structure of the beam described in Section 4.3.

The electric field inside the cavities oscillates at a 400 MHz frequency, such that after reaching the maximum energy, an ideally timed proton (with exactly the desired

energy) does not experience any accelerating force while passing through the cavities. However, if a proton has a slightly lower (higher) energy, it gets accelerated (decelerated) towards the ideal timing.

The high oscillation frequency is obtained with *klystron* tubes, which contain electron beams and are placed next to the cavities. The electron beams are modulated at 400 MHz, and the power is transmitted to the cavities via metallic *waveguides*. To minimize losses, also the radio frequency cavities are operated in the superconducting state. For this purpose, they are placed inside four cylindrical *cryomodules* that cool them down to 4.5 K.

4.3 Fill structure and luminosity

The LHC beam is not a continuum of protons, but consists of thousands of *bunches* of protons, which are typically $\approx$10 cm long. For the LHC Run 2, the *bunch spacing*, i.e. the minimum distance between two subsequent bunches, was decreased from 50 to 25 ns, corresponding to a bunch separation of 7.48 m. The bunch spacing also determines the bunch crossing frequency, which is 40 MHz for 25 ns bunch spacing.

The number of bunches in a fill (n_b) and the number of protons per bunch (N_p) depend on the specific *filling scheme* used when injecting the beam from the SPS to the LHC. The filling scheme defines which of the 26659 m/7.48 m = 3564 *buckets* in the LHC beam are filled with a proton bunch and which are left empty. In typical filling schemes, the bunches are arranged in *bunch trains* of filled buckets, separated by series of empty buckets. The empty buckets are needed to ensure correct injection of the beam from the SPS to the LHC, as well as for controlled beam dump, as the *kicker magnets* used to direct the beam in and out of the LHC need time to ramp up their magnetic fields. The maximum n_b that can be achieved with the current LHC is 2808, but in 2016, when the data used in this thesis was collected, the most commonly used filling scheme had 2208 bunches. In this scheme, the bunches are organized in 46 trains of 48 bunches per train, with each bunch containing approximately 11.5×10^{11} protons at the start of the fill.

The production rate dN/dt for a specific process is determined by its cross section σ and the *instantaneous luminosity* $\mathcal{L}$:

$$\frac{dN}{dt} = \sigma \mathcal{L}. \tag{4.2}$$

Figure 4.3: The peak instantaneous luminosity measured at the interaction point of the CMS experiment from 2010 to 2018. [109]

Integrating over a time period T, we obtain the *integrated luminosity L* which is proportional to N and thus reflects the amount of data collected:

$$N = \sigma \int_0^T \mathcal{L} dt = \sigma L. \tag{4.3}$$

The instantaneous luminosity is typically measured in units of $[\,\text{cm}^{-2}\,\text{s}^{-1}\,]$, or alternatively in $[\text{Hz}/\,\text{nb}] = 10^{-33} \times [\,\text{cm}^{-2}\,\text{s}^{-1}\,]$. The maximum instantaneous luminosities per day for the LHC Run 1 and Run 2 are shown in Figure 4.3. It shows that in 2016 the LHC reached for the first time its design instantaneous luminosity of $10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$, and then surpassed it by a factor of 1.5. Later in 2017–2018, even two times the design luminosity was achieved. This is the maximum value currently allowed by the inner triplet magnets, which get heated by debris from the collisions at high luminosities. The resulting integrated luminosities for each operation year during Run 1 and Run 2 are summarized in Figure 4.4.

The different factors that determine the instantaneous luminosity can be summarized in the following equation:

$$\mathcal{L} = \frac{N_{\text{p}}^2 n_b f_{\text{rev}} \gamma}{4\pi \epsilon_n \beta^*} F. \tag{4.4}$$

Table 4.1: The LHC beam parameters for 13 TeV operation. The nominal parameters are compared to typical values used in 2016.

Parameter	Design value	Typical value in 2016
n_b	2808	2208
N_p	1.15	1.15
f_{rev} (kHz)	11.25	11.25
γ	7461	7461
ϵ_n (μm)	3.75	2.20
β^* (at CMS, cm)	55	40
F	0.48	0.65
$\sqrt{s}$ (TeV)	14	13
$\mathcal{L}_{peak}$ (10^{34} cm^{-2})	1.0	1.4

Here f_{rev} is the bunch revolution frequency and γ is the relativistic Lorentz factor, both of which are determined by the beam energy. The *transverse emittance* ϵ_T measures the spread of the beam particle positions and momenta in the transverse direction. The *normalized transverse emittance* present in Eq. (4.4) is defined as $\epsilon_n = \beta\gamma\epsilon_T$ (where β and γ are standard relativistic functions), which does not depend on the beam energy. The symbol β^* refers to the value of the *betatron function*, which describes the focusing of the beams at the interaction point, and is determined by the properties of the focusing magnets and the crossing angle of the beams. Finally, the *form factor F* gives the geometrical suppression in luminosity due to non-zero crossing angle at the interaction point. The design values for these parameters, and the values typically used during 2016 data taking, are listed in Table 4.1.

4.3.1 Pileup

According to Eq. (4.2), the instantaneous luminosity determines the interaction rate at a given $\sqrt{s}$, for which the cross section is constant. Therefore at high luminosities it is possible to achieve an interaction rate that is larger than the bunch crossing rate, meaning that on average several pp interactions occur per bunch crossing. This phenomenon is called *pile-up*, and usually it is quantified in terms of the *mean number of interactions per bunch crossing* $\langle N_{PU} \rangle$.

Figure 4.4: The overall integrated luminosity measured for the data collected by the CMS experiment from 2010 to 2018. [109]

The amount of pileup can be calculated from the instantaneous luminosity by multiplying it with the total inelastic pp cross section $\sigma_{\text{in}}^{\text{pp}}$:

$$\langle N_{\text{PU}} \rangle = \sigma_{\text{in}}^{\text{pp}} \mathcal{L}. \tag{4.5}$$

The distributions of $\langle N_{\text{PU}} \rangle$ for different years of LHC operation are shown in Figure 4.5. The values of $\sigma_{\text{in}}^{\text{pp}}$ used for different center-of-mass energies are also listed.

The additional pp interactions that occur in the same bunch crossing as the hard collision are referred to as *in-time pileup*. In addition, the decay times of calorimeter signals are typically longer than 25 ns (one bunch crossing), so the signal readout requires that the measured pulse is decomposed into in-time and out-of-time components. Incorrect pulse decomposition, as well as incorrect timing of detector signals (from calorimeters as well as faster subdetectors) can cause mixing of signatures from consecutive bunch crossings. Therefore interactions taking place in the preceding or subsequent bunch crossings can mix with the signals from the hard event, manifesting as *out-of-time pileup* in the reconstructed events.

As the total inelastic pp cross section is dominated by soft QCD interactions, they constitute most of the pileup interactions. Even though these interactions only produce low-energy particles, they still deposit energy in the detector which needs to be taken into account so that they do not bias the measurement of the hard collision.
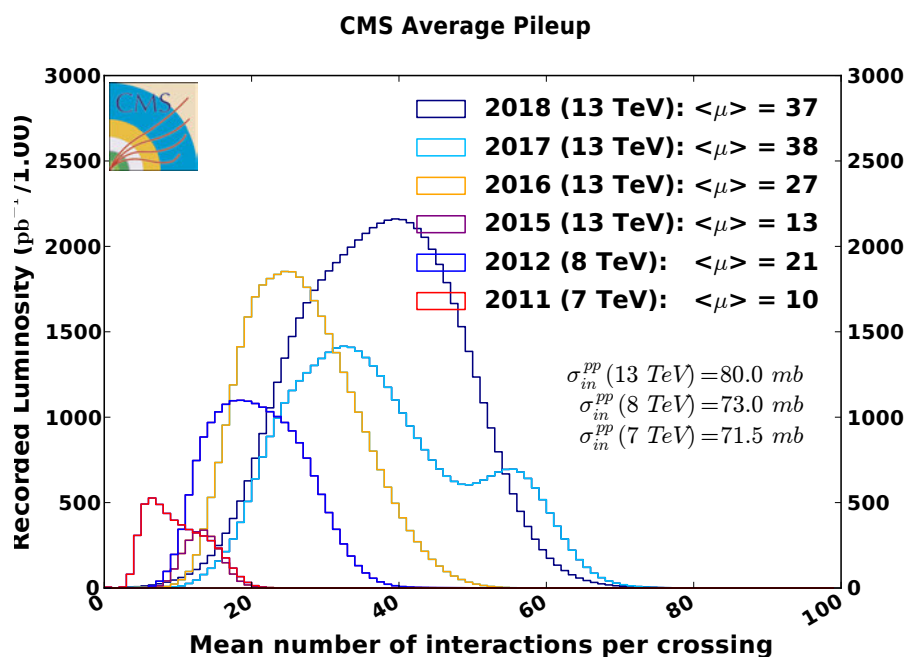
Figure 4.5: The pileup distribution observed at the CMS experiment, expressed as the mean number of interactions per bunch crossing, from 2010 to 2018. The distributions are normalized to the integrated luminosity of the data collected each year. [109]

The increase in luminosity from Run 1 to Run 2, combined with the shortening of the bunch crossing time from 50 ns to 25 ns, mean that the pileup observed in the CMS experiment has been larger during Run 2 compared to Run 1, as can be seen in Figure 4.5. The increased pileup is one of the main challenges in the analysis of the Run 2 data. Thus we will return to this topic multiple times in the following chapters: the upgrade of the Level-1 trigger, presented in Chapter 6 was largely motivated by the pileup increase. The pileup also needs to be included in the simulation of the collision events as described in Chapter 7, and mitigated in the reconstruction of collisions as discussed in Chapter 8.

Chapter 5

Compact Muon Solenoid

5.1 Overview

The Compact Muon Solenoid (CMS) experiment is located at the LHC access point 5, in the village of Cessy in France. The detector is situated 100 meters underground, around an *interaction point* where the beams of the LHC are set to collide. While CMS is the heaviest of the LHC detectors, weighting 14000 tonnes, its dimensions of $22 \times 15\,\text{m}$ make it relatively *compact* compared for example to the ATLAS experiment ($46 \times 25\,\text{m}$). Precise, high-efficiency *muon* measurement, enabled by large gas chamber detectors, is one of the distinctive properties of the CMS experiment. The detector is constructed inside and around a large superconducting *solenoid* magnet, which is $13\,\text{m}$ long and $6\,\text{m}$ in diameter. It provides a strong magnetic field of $3.8\,\text{T}$ in the inner part of the detector, bending the trajectories of charged particles for precise measurement of momentum and charge. The detector is designed to measure proton-proton collisions, as well as heavy ion collisions (e.g. lead-lead and proton-lead interactions). Herein we focus on the proton-proton collisions. In the following, the detector is described in the configuration it had in 2016, when the data used in this thesis were recorded. The upgrades performed after 2016 and to be carried out in the future are summarized in Section 5.6.

Figure 5.1 illustrates the structure of the CMS detector. The protons are set to collide in the *interaction point* at the geometrical center of the detector. After traversing a thin ($< 1\,\text{mm}$ thick near the interaction point) beryllium beam pipe, the particles produced in the interactions encounter several layers of different *subdetectors* which surround the interaction point, forming an onion-like structure. They are used to measure the trajectories and energies of the particles produced in collisions. The detector is built

Figure 5.1: An opened view of the CMS detector and its main subsystems. [110]

with cylindrical geometry: in the barrel part, the detector layers are situated around the beam axis, while their counterparts in the endcaps are circular disks.

The subdetector closest to the interaction point is the *silicon tracker*, which measures trajectories of charged particles in the magnetic field. The tracker is surrounded by an *electromagnetic calorimeter* (ECAL), designed to absorb electrons and photons and measure their energies. In the endcaps, the ECAL is complemented by separate *preshower* detectors. The charged and neutral hadrons can initiate an electromagnetic shower in the ECAL. The next subdetector layer known as the *hadronic calorimeter* (HCAL) absorbs the rest of the shower energy, as well as any hadronic showers initiated in the HCAL itself. In order to have a detector as hermetic as possible, the forward region close to the beam pipe is equipped with additional *forward calorimeters*.

To minimize the energy loss of the particles before their energies are measured in the calorimeters, the superconducting coils are placed outside the calorimeters. After the coils, several layers of steel return yoke direct the magnetic flux to keep the strong magnetic field contained inside the detector volume. They also provide mechanical

support for the whole construction. As the only charged particles to traverse the detector all the way to the return yoke are the muons, the large detector volume outside the solenoid is filled with *muon chambers*, interleaved between the steel layers. They are gaseous detectors that provide high-efficiency muon identification and high-resolution momentum measurement.

In the CMS experiment, a right-handed coordinate system, illustrated in Figure 5.2, is used. The origin corresponds to the nominal interaction point at the center of the detector. The *x* axis is defined to point towards the center of the LHC ring, while the *y* axis points up towards the sky, perpendicular to the plane defined by the LHC ring. The *x* and *y* axes span the *transverse* plane, where the azimuthal angle ϕ is defined. The *z* axis corresponds to the longitudinal axis of the CMS cylinder and points along the direction of the anticlockwise beam (as seen from above). The polar angle θ is measured with respect to the positive *z* axis.

In practice, the *pseudorapidity* η is commonly used instead of θ. It is defined as

$$\eta = -\ln(\tan\frac{\theta}{2}). \tag{5.1}$$

Figure 5.2: Illustration of the CMS coordinate system. While traditionally the *z* axis corresponds to the *vertical* coordinate (as drawn here), in the CMS coordinates it corresponds to the *horizontal* direction of the beam.

For highly relativistic particles with $m \ll p$, η is equal to the *rapidity* defined as

$$y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z}. \tag{5.2}$$

The CMS detector has approximately 100 million readout channels, managed by a complex data acquisition system. It relies on a two-level trigger system, presented in detail in Chapter 6. In the following, an overview of the main subsystems of the CMS is provided. A more detailed description of the CMS detector can be found in Ref. [111].

5.2 Tracking system

The silicon tracker is designed to measure trajectories of charged particles in three spatial dimensions with a high resolution. The charge and momentum of a particle can be deduced from the curvature of its trajectory. The structure of the tracker system is illustrated in Figure 5.3. The system is cylinder-shaped, with a diameter of 2.6 m and a length of 5.8 m. The part closest to the interaction point is equipped with a fine-granularity *pixel detector*, while the large bulk of the tracker region consists of larger *silicon strips*.

Until the end of 2016, the pixel detector contained three cylindrical layers in the barrel and two discs in each endcap, with 66 million pixels in total, providing a resolution of $150 \times 100 \, \mu\text{m}^2$ per pixel. The closest parts of the pixel detector were situated only 4.4 cm away from the nominal interaction point. The configuration of layers and discs is illustrated in the lower part of Figure 5.4. The upper part shows the upgraded pixel detector discussed further in Section 5.6.

As shown in Figure 5.3, the majority of the tracker volume is occupied by a *strip tracker* that contains in total 9.3 million silicon strips, each 80–180 μm wide. The silicon strips are arranged in a configuration visible in Figure 5.3, with four strip layers in the inner barrel (TIB), six layers in the outer barrel (TOB), three layers in the inner disks (TID) and nine layers in each endcap (TEC- and TEC+). With this configuration, a spatial resolution of 10–100 (20–1000) μm in the transverse (longitudinal) direction is achieved.

Most of the strip modules in the barrel are aligned with the z axis, i.e. parallel to the beam, providing information on r (the radial coordinate) and ϕ (the azimuthal angle).

Figure 5.3: A schematic picture of the CMS tracker system. The high-resolution pixel detector is surrounded by a large strip silicon tracker, which consists of inner barrel (TIB), outer barrel (TOB), inner disks (TID) and endcaps (TEC- and TEC+). The double lines indicate module pairs providing three spatial coordinates. [111]

In the endcaps, they are aligned radially, providing information on z and ϕ. Some of the modules, however, are tilted by a *stereo angle* of $100 \, \text{mrad}$ ($\approx 5.7°$), and attached back-to-back with another module, hence providing the third spatial coordinate. These stereo modules are marked as double lines in Figure 5.3.

The pixel and silicon trackers are both based on ionization of doped and layered semiconductor silicon with a p–n junction. A reverse bias voltage is applied between arrays of pixels or strips, creating a *depletion region* with no free charges. As a charged particle traverses the semiconductor, thousands of electron-hole pairs are created by ionization. The electrons drift in the electric field created by the bias voltage, until they reach an electrode where the charge is collected and read out.

A *radiation length* is defined as the distance where an electron loses $1 - 1/e = 63\%$ of its energy. The amount of material in the tracker corresponds to approximately 1.8 radiation lengths. Therefore electrons and photons have a high probability to interact with the tracker material via bremsstrahlung and electron conversions, and this needs to be taken into account in the interpretation of the tracker signals. Heavier particles, such as hadrons and muons, are *minimum ionizing particles*, which typically traverse the tracker almost without energy losses. Thus their energies can be accurately measured once they reach the outer detector layers.

Figure 5.4: Comparison of the original CMS pixel detector used until the end of 2016 (lower half of the image) and the new pixel detector installed in early 2017 (upper half of the image). [112]

The tracker system provides excellent momentum resolution of $< 1\%$ for charged hadrons with $p_T < 20\,\text{GeV}$. The resolution decreases with higher momenta, reaching $\approx 5\%$ at $1\,\text{TeV}$. As high-energy particles are likely to reach calorimeters or muon chambers, the measurement of their momenta in the tracker can be complemented by information from other subdetectors, as discussed later in Chapter 8. The tracker information is used to reconstruct tracks of charged particles (Section 8.1.1) and vertices (Section

5.3 Electromagnetic calorimeter

The silicon tracker is surrounded by the *electromagnetic calorimeter* (ECAL). The ECAL is designed to absorb electrons and photons and measure their energies. It is built of a single layer of lead tungstate ($PbWO_4$) crystals.

Lead tungstate has an exceptionally high density of $8.28\,\text{g/cm}^3$ (higher than e.g. the density of steel). Due to this density and the high atomic numbers of lead and tungsten, a single layer of 22–23 cm long crystals is enough to achieve ≈ 25 radiation lengths, meaning that electrons and photons have a very high probability to interact with the crystal material. In the vicinity of a heavy nucleus, electrons are likely to radiate bremsstrahlung photons, which in turn can convert into electron-positron pairs. When these processes follow each other repeatedly, a collimated *electromagnetic shower* is

created. According to simulations, in total $\approx 98\%$ of the e/γ energy (up to 1 TeV) is absorbed in the ECAL, and only the remaining small fraction ends up in the HCAL.

The lead tungstate also acts as a scintillator, emitting low-energy photons as the excited atoms return to lower energy states. The scintillation light is detected with photodiodes connected at the end of the crystals. In the absence of radiation damage, the $PbWO_4$ crystals are optically as transparent as glass, and the scintillation light can reach photodiodes almost without losses.

Lead tungstate also has a short excitation time: approximately 80% of the scintillation light is emitted within 25 ns (time of one bunch crossing). Avalanche photodiodes (APD) collect the scintillation light from the crystals. The amount of light measured in the photodiodes is proportional to the energy deposited by the original particle that initiated the electromagnetic shower.

The *hadronic interaction length* (simply "interaction length" from hereon) corresponds to an average distance that a hadron can travel in the material before a nuclear interaction occurs. The material in the ECAL corresponds to one interaction length, so most hadrons start showering already inside the ECAL, depositing their energy partially in the ECAL and partially in the HCAL.

Due to its simple structure and high density, the ECAL is compact enough to fit inside the solenoid magnet. This type of calorimeters, where the same material acts as absorber and scintillator, are referred to as *homogeneous calorimeters*.

The ECAL is a fine-grained detector. It has in total 75 848 calorimeter crystals, each of which produces a separate signal. The barrel part of the detector covers the region with $|\eta| < 1.479$. It is segmented into 360 crystals in the ϕ direction, and 170 crystals in the z direction, summing up to 61200 crystals in total. The ECAL crystals are grouped into 36 *supermodules*, each corresponding to 10 degrees in ϕ. The supermodules contain 4 *modules*, each with separate readout. The configuration of crystals and modules is illustrated in Figure 5.5.

In the barrel, each crystal is 23 cm long, with square-shaped front face of $2.2 \times 2.2\,cm^2$ and slightly larger rear face of $2.6 \times 2.6\,cm^2$. This means segmentation of approximately 0.02×0.02 in the (η, ϕ) plane, which roughly corresponds to the transverse size of electromagnetic showers in $PbWO_4$. The crystals do not point directly to the nominal interaction point. Instead, they are tilted away by 3 degrees to avoid situations where a particle would traverse the ECAL through a small gap between two crystals.

Figure 5.5: Illustration of the structure of the ECAL. The PbWO$_4$ crystals are grouped into modules and supermodules in the barrel, and into supercrystals in the endcaps. The preshower detectors provide additional sensitivity in the endcaps. [111]

The ECAL endcaps are situated 3.15 m away from the interaction point and both of them contain 7324 crystals placed in a rectangular grid. The endcaps extend the $|\eta|$ coverage of the ECAL up to 3.0. The endcap crystals are 23 cm long, with front (rear) face of $2.9 \times 2.9\,\mathrm{cm}^2$ ($3.0 \times 3.0\,\mathrm{cm}^2$). Vacuum phototriodes (VPT) are used to collect the scintillation light. They are set to point 13 cm beyond the nominal interaction point, effectively tilting them by 2–8° with respect to the interaction point. In the endcaps, blocks of 5×5 crystals known as *supercrystals* share a common readout.

In the endcaps, fine-grained two-layer *preshower* detectors are installed in front of the ECAL endcap disks. They cover a pseudorapidity range of $1.65 < |\eta| < 2.6$. The preshower detectors are *sampling calorimeters* with separate *passive* layers for absorption (initiation of the electromagnetic shower) and *active* layers for the detection of the shower. Both of them contain two disks, separated by 20 cm from each other. Each disk contains a lead absorber layer where electromagnetic showers are created and a silicon strip sensor layer that measures the energy of the showers. The strips in the two layers are orthogonal to each other, separated from each other by a transverse distance of 1.9 mm. The silicon strips are only 2 mm wide, so the spatial resolution is significantly higher than obtained by the ≈3 cm wide ECAL endcap crystals. This high

resolution is utilized especially to identify neutral pions. As photon pairs from π^0 decay can be nearly collinear in the endcap, they could be mistaken for one energetic photon without the additional resolution provided by the preshower detector.

5.4 Hadron calorimeter

Hadrons (neutral or charged) are usually not fully absorbed in the ECAL. After traversing the ECAL, they encounter the *hadronic calorimeter* (HCAL). It is designed to stop all particles that reach it and measure their energy. The only exceptions are muons, which pass through the HCAL all the way to muon chambers, and neutrinos or other weakly interacting particles that escape the whole CMS detector without leaving a trace.

The HCAL is a sampling calorimeter, with alternating passive layers of brass absorber and active layers of plastic scintillator material. In the passive layers the hadrons collide with the nuclei in the brass medium, creating cascades of secondary particles. When this *hadronic shower* reaches an active layer, scintillation light is created and measured by photodetectors. As part of the energy is deposited in the passive absorber material, where it cannot be measured, the total energy deposited in the HCAL is estimated indirectly, based on the information from the active layers. As discussed above, typically the hadrons start showering already in the ECAL, so the signals from ECAL and HCAL need to be combined to obtain a correct estimate of the hadron energy.

The structure of the hadronic calorimeter system is shown in Figure 5.6. The HCAL barrel (HB) covers the pseudorapidity region of $|\eta| < 1.3$. It is divided into 36 azimuthal *wedges*, each 1 m thick. The HB contains 14 brass plates parallel to the z axis, each 51–57 mm thick. In addition, 40 (75) mm thick steel plates are placed in the front (back) face of the system. The plastic scintillator layers contain 16 η segments with separate readouts. In the barrel, the amount of material in the HCAL absorbers corresponds to 6–11 interaction lengths, such that the amount of material traversed by hadrons increases with $|\eta|$.

The HB is complemented by an outer HCAL layer (HO), also known as the *tail catcher*, placed outside the solenoid coil. It contains one layer of absorber iron equipped with scintillators, placed near the innermost layer of the return yoke.
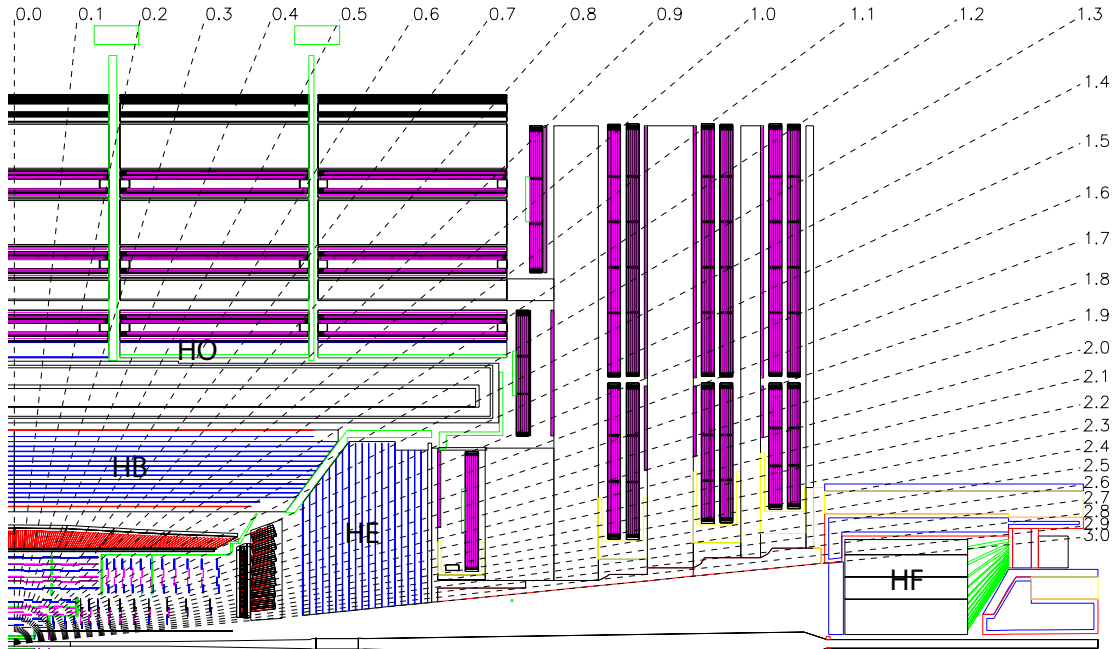
Figure 5.6: Illustration of one quadrant of the hadronic calorimeter system, shown in (z, y) plane. The HCAL barrel (HB) and endcap (HE) detectors are completed by the outer barrel (HO) and forward (HF) calorimeters. The muon systems are shown in purple. [111]

The HCAL endcaps (HE) cover the pseudorapidity range of $1.3 < |\eta| < 3.0$. The HB and HE are inside the solenoid coil, so the trajectories of charged particles are bent in the magnetic field all the way to HCAL.

In the barrel, the HB provides an $\eta \times \phi$ granularity of approximately 0.087×0.087. The granularity becomes coarser with increasing $|\eta|$, reaching approximately 0.17×0.17 in the endcaps. Even though the segmentation of the HCAL is about 25 times coarser than that of the ECAL, it is still fine enough to separate neutral hadron energy deposits from those originating from charged hadrons bent by the magnetic field.

With all calorimeter layers in the ECAL and HCAL combined, the amount of material corresponds to approximately 12 (10) interaction lengths in the barrel (endcaps). Therefore the probability for punch-through, i.e. for hadrons to pass through the HCAL and hit the muon chambers, is very small, although not zero.

In order to make the calorimeter system as hermetic as possible, the HCAL is complemented by *hadron forward calorimeters* (HF), situated ± 11.2 m away from the nominal
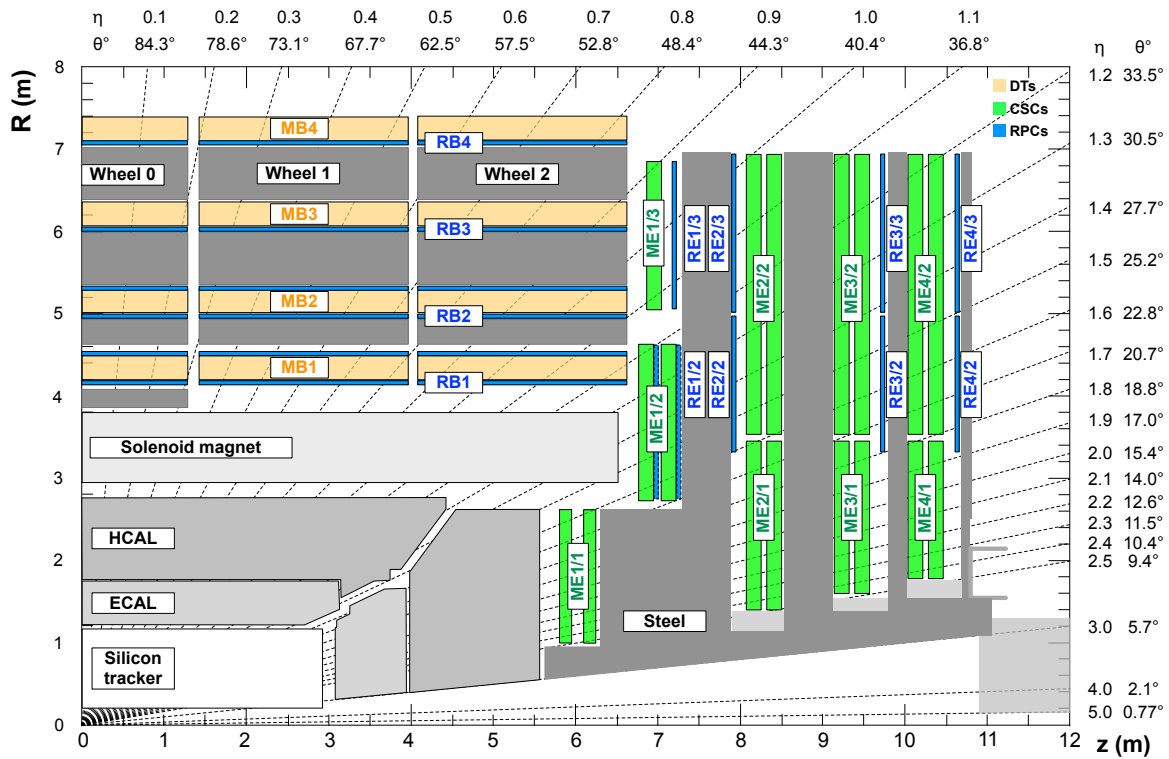
Figure 5.7: The CMS muon chamber configuration illustrated in one quadrant of the detector. The muon stations in the barrel (MB) are drift tubes (DTs), while in the endcaps (ME) they are cathode strip chambers (CSCs). The resistive plate chambers (RPC) provide redundancy in the barrel (RB) and endcaps (RE). [113]

interaction point. They extend the $|\eta|$ coverage to up to 5.2. The HF detectors are cylindrical structures with a diameter of 2.6 m. Due to intense radiation conditions close to the beam, they are built of steel absorber, with quartz fibers as the active elements, read out by photomultipliers. The $\eta \times \phi$ resolution of the HF is 0.175×0.175, and the amount of material corresponds to $\approx$10 interaction lengths.

The CMS calorimeter system is of great importance in precise measurements of jets, as well as neutrinos or other particles manifesting as missing transverse momentum. The HF is also used for measuring the instantaneous luminosity. The calibration and combination of signals from ECAL and HCAL are discussed later in Section 8.1.

5.5 Muon chambers

As shown in Figure 5.1, the majority of the CMS detector volume is occupied by *muon chambers*, interleaved between the steel plates of the magnetic return yoke.

The CMS muon system contains three types of chambers, all of which use gas as active medium. The chambers are arranged in layers known as *muon stations*. The barrel part ($|\eta| < 1.0$) contains four layers of aluminum *drift tubes* (DTs), while in the endcaps ($0.9 < |\eta| > 2.4$), four layers of *cathode strip chambers* (CSCs) are used. To improve efficiency especially in the transition regions between the barrel and the endcaps, four layers of *resistive plate chambers* (RPCs) complement the DTs and CSCs up to $|\eta| = 1.7$.

The DTs are detector elements filled with ionizable gas (85/15% mixture of argon and CO_2), equipped with a positively-charged wire spanned through the center of each tube. When a muon traverses a tube, it ionizes the gas, creating secondary electrons that drift towards the wire in the electric field that surrounds it. As the drifting electrons are accelerated by the field, they ionize the gas further, creating an avalanche of secondary electrons. Once the electrons reach the wire, they create an electric signal. The drift time is typically a few hundred nanoseconds, depending on the distance between the muon and the wire.

In total, the CMS muon system contains 250 DT chambers placed in five barrel wheels. Part of them are installed with wires parallel to the z axis, so that they measure the r and ϕ coordinates, while others are installed orthogonally, measuring the r and z coordinates. A typical *single-hit* resolution, i.e. the resolution for the estimated distance between the muon track and the anode wire, is 200-300 μm (300-800 μm) for the chambers placed parallel (orthogonal) to the z axis.

The CSCs work in a similar way as DTs, except that they contain a large number of electrodes that collect the signals. The CSCs are gas containers equipped with alternating layers of negatively-charged cathode strips and positively-charged anode wires. The cathode strips are placed radially, providing precise measurements of the r and ϕ coordinates, whereas the wires provide a coarse estimate of the z coordinate.

In total, there are 540 CSC chambers in the CMS endcaps. Each chamber has six anode wires placed between cathodes. The radial anode wires and longitudinal cathode strips are orthogonal to each other. The CSCs use a 40/50/10% mixture of argon, CO_2 and CF_4 as the active gas. Due to a larger density of electrodes, the CSCs provide better spatial resolution than DTs. In the CSCs typical spatial resolution for reconstructed hits is $\approx 100\,\mu$m.

In the RPCs, the gaseous volume is sandwiched between two resistive plates. The outer surfaces of the plates are coated with conductive graphite layers to apply voltage across the gaseous volume. When an avalanche of electrons is created in ionization

of gas, it is collected by metallic strips connected to the anode plate. The RPCs use a 95.2/4.5/0.3% mixture of freon ($C_2H_2F_4$), isobutane (C_4H_10) and sulfur hexafluoride (SF_6) as the active gas. In total, 1056 RPC chambers are placed in four layers both in the barrel and in the endcaps. They have a prompt time response of $\approx 2\,$ns and spatial resolution of $\approx 1\,$cm.

5.6 Upgrades

To keep up with the increasing LHC luminosity, and to benefit from technological advances in instrumentation and electronics, also larger upgrades of the subsystems are performed. The CMS *Phase-1 upgrade* program started after the end of Run 1 and will continue until 2021, when LHC Run 3 starts. It includes a complete upgrade of the Level-1 trigger system [114], HCAL and HF photodetector and readout upgrades for additional granularity [115], and installation of a new pixel detector [116].

The upgraded Level-1 trigger system is described later in Section 6.1. In the HCAL (HF), the old photodiodes are replaced by silicon photomultipliers (multi-anode photomultipliers), which are expected to perform better under high-luminosity conditions. Also the front-end electronics are replaced, increasing the number of readout channels.

As part of the Phase-1 upgrade, the whole pixel detector was replaced with a new one in early 2017. The new pixel detector contains four barrel layers and three discs in the endcaps, doubling the number of pixels to 123 million [112]. The innermost layer in the barrel with a radius of 3.0 cm is even closer to the beam than previously. The difference between the old and new configurations is illustrated in Figure 5.4. The upgrade was needed because the old detector was suffering from serious radiation damage and from dynamic inefficiency under high-pileup conditions due to the limitations of the readout system. In the new detector, the amount of passive material is substantially reduced by placing the upgraded readout system further away from the detector, and with other measures such as a new cooling system design. The new system is more robust against high pileup due to the new pixel layers, and a new faster readout system with an increased number of channels.

After the end of Run 3, the next major upgrade of the detector will take place. The goal of the *Phase-2 upgrade* is to prepare for the intense luminosity conditions of the high-luminosity LHC, expected to start operating after the third long shutdown of the LHC, which will take place in 2024–2026. The plans for Phase-2 upgrade are documented in Ref. [117].

Chapter 6

Online event selection

As discussed in Chapter 5, the CMS detector has approximately 100 million readout channels in total. When all signals associated to a single event (measured to originate from one 25 ns bunch crossing) are read out, filtered, shaped, compressed and zero suppressed by the data acquisition (DAQ) system, they result in approximately one megabyte of data per event. The 25 ns bunch crossing period of the LHC corresponds to 40 MHz nominal bunch crossing rate. The CMS DAQ system is *synchronous*, meaning that the signals are sent to buffers and processed by the DAQ system at this rate. Thus if data from every bunch crossing would be permanently stored, the experiment would produce 40 terabytes of data per second, equal to 144 petabytes per hour. Reading out and storing of such amount of data is simply impossible with the existing technology—currently, all the data ever collected at CERN experiments over decades, stored in large data centers around the world, sum up to a few hundred petabytes. On the other hand, as discussed in Chapter 4, only a tiny fraction of these events contain interesting physics such as production of heavy particles.

The solution is to equip the detector with a *trigger* system, which is designed to monitor the collisions in real time and identify the events that probably contain interesting physics. At every bunch crossing, the trigger system decides whether the event is *accepted* (read out and stored), or discarded. The concept of triggering was introduced in particle physics already a century ago, when it was noted that "visual or audible methods of counting are quite trying on the nerves", so "a self-recording device would therefore be an obvious improvement" [118]. Triggering is also commonly referred to as *online selection* of events, as opposed to *offline selection* where a subset of recorded and stored events is selected for further analysis.
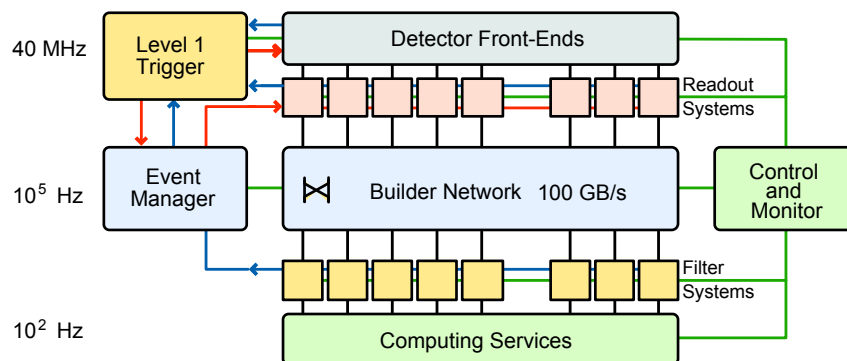
Figure 6.1: An overview of the CMS trigger and DAQ system. The hardware-based Level-1 trigger receives data from detector frontends and readout systems (blue lines) and selects events to be fully read out from the detector, for which it sends a "Level-1 Accept" signal (red lines). After the full information from all subdetectors is read out (black lines) and combined by a fast event builder network system, a software-based High-Level Trigger ("Filter systems") select the events to be permanently stored. Green lines indicate control and monitoring data flows. [111]

The trigger system defines the physics reach of an experiment, because its settings limit the types of events which are stored for further analysis, and hence the range of possible offline analyses. Designing a reliable trigger system and ensuring its correct functioning are a crucial part of the operation of a collider experiment, because once an event is discarded by the trigger, it is permanently lost and cannot be recovered.

In modern collider experiments, the trigger system is typically composed of several *levels*, such that each level discards part of the events and sends the remaining ones to the next level for closer inspection. This approach allows algorithms of increasing complexity to be applied at each level. The CMS trigger system contains two levels: the *Level-1 trigger* (L1T) selects the events that will be fully read out. Its maximum output event rate, i.e. the *trigger rate* of $\approx 100\,\text{kHz}$ is determined by the bandwidth of the readout system. The buffer sizes in the DAQ system are limited, so the L1T has only a few microseconds to make a decision before the buffers start to overflow, leading to uncontrolled loss of data known as *dead time*. During Run 2, the L1T latency was measured to be $3.8\,\mu$s. To achieve this low latency, the L1T is completely hardware-based. Once the events accepted by the L1T are fully read out, they are sent to the *High-level trigger* (HLT), which is a software-based filtering system running on a large computer farm. The HLT reconstructs the events using a streamlined version of the standard CMS offline event reconstruction software, reducing the event rate under

1 kHz, allowed by the storage system. A diagram of the trigger and DAQ systems in the CMS is shown in Figure 6.1. Since the author contributed actively to the quality control of the L1T system during Run 2, the L1T is presented in detail in Section 6.1, followed by a short discussion of the HLT in Section 6.2.

6.1 Level-1 trigger

To achieve low latency in trigger decisions, the L1T is built of customized hardware processing units that perform low-level logic operations based on input from different CMS subsystems. The processing of signals and the *trigger algorithms* that evaluate each event are implemented in field programmable gate arrays (FPGAs). The workhorse of the current CMS L1T system are large FPGAs of Xilinx Virtex 7 family, which are used to perform logical operations in all processor boards of the trigger system. The data are transferred via fast serial optical links that transmit up to 10 gigabits per second. The architecture of the whole trigger system is modular, and designed to be easily expandable, in case changing LHC running conditions or improved trigger algorithms require additional logic or memory resources.

Since reading out the full event data at 40 MHz rate would be too bandwidth-consuming, the L1T takes input only from selected subsystems: ECAL, HCAL (including HF) and the muon chambers. As even the full information from these subdetectors would take too long to process, the front-end electronics of the detector elements perform partial averaging of the information to build simpler *trigger primitives* that the L1T can take as input. The detailed and therefore abundant tracker information is not used at all by the L1T.

Figure 6.2 shows the architecture of the CMS Level-1 trigger in a schematic way. The muon trigger, presented in Section 6.1.3, takes input from muon chambers and runs the muon-related trigger algorithms. The calorimeter trigger, discussed in Section 6.1.4, runs several algorithms targeting different types of calorimetric objects. The final "L1Accept" (L1A) decision to read out an event and send it to HLT is performed by the global trigger as described in Section 6.1.5. The efforts to ensure the correct functioning of the Level-1 trigger and the procedure used to certify the data collected during Run 2 are described in Section 6.1.7.

In 2013–2017, the L1T system was completely upgraded to ensure good trigger efficiency and performance under the harsh luminosity conditions of Run 2. The upgrade was crucial for the success of the whole CMS physics program, since a factor of six
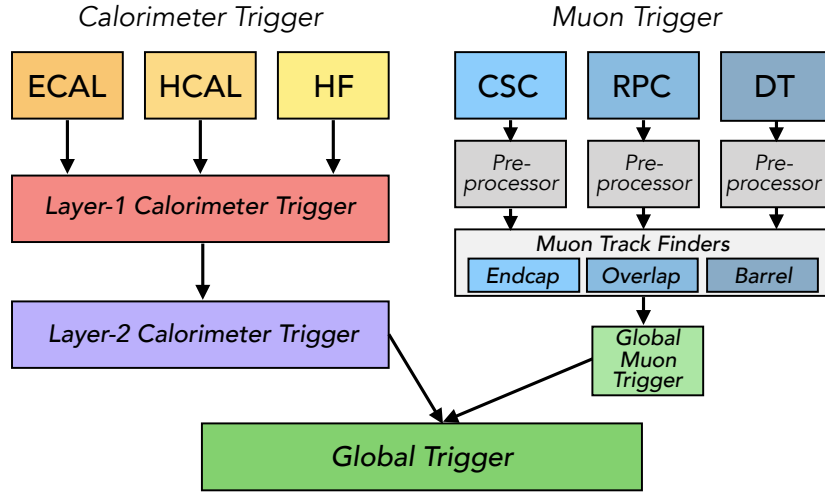
Figure 6.2: The architecture of the CMS Level-1 trigger system in Run 2. The purpose of each module is explained in the text.

increase in trigger rates was predicted if the Run 1 trigger were to be used in Run 2. The insupportable trigger rates would force the operators of the experiment to discard most of the interesting events by applying higher trigger thresholds or by other means, which would be detrimental for the physics goals of the experiment. The *Phase-1 upgrade* covered the hardware, the firmware, the timing control system, and the trigger control software. Also the trigger algorithms were redesigned to fully benefit from the upgraded architecture. The calorimeter trigger was upgraded during the long shutdown 1 (2013–2015), and the muon trigger in 2016–2017. Here, the main features of the upgraded version of the L1T is described. A more detailed description is provided in Ref. [114]. The original L1T system used during Run 1 is described in Ref. [119]. The next major trigger upgrade, the *Phase-2 upgrade*, is foreseen after Run 3, which is planned to end in 2023. The plans for Phase-2 upgrade are documented in Refs. [117, 120].

6.1.1 Trigger algorithms

The L1T runs simultaneously a large number of different *trigger algorithms* designed to identify specific types of collision events. Typically the trigger algorithms look for exceptionally large localized energy deposits in the detector elements, corresponding to high-p_T particles. From the energy deposits, they aim to reconstruct the highest-p_T particles in an approximate way, building *trigger objects* corresponding to muons, jets, hadronically decaying tau leptons, photons or electrons (e/γ objects), and different

energy sums in the transverse plane, such as the scalar jet p_T sum (H_T), the missing transverse momentum (p_T^{miss}), and its jet-based equivalent called H_T^{miss}. The properties of the trigger objects are then compared to predefined *trigger thresholds*, which are typically defined in terms of p_T. A Level-1 trigger algorithm with a specific threshold (or a combination of thresholds or multi-object algorithms) is referred to as a *trigger seed*.

Low threshold(s) for a seed mean that more events are accepted, so the seed produces a higher trigger rate. The final decision to accept or discard an event is made in the *Global trigger* as a logical OR of all the seeds which are used simultaneously. Thus if any of the seeds accepts the event (and if some additional basic conditions, such as presence of a proton bunch crossing, are met), the event is accepted by the L1T. In this case, the trigger system sends a "Level-1 Accept" (L1A) signal to the DAQ system, allowing it to proceed to read out the full event from the buffers and send it to the HLT.

While the details of different trigger algorithms vary a lot, most of them share some common principles and design goals. Typically, a trigger algorithm looks for high-p_T physics objects inside a geometrical region limited by a maximum value of $|\eta|$. This is an efficient way to select hard events with high momentum transfer, while rejecting most of the abundant soft QCD interactions that produce low-p_T particles with high $|\eta|$. Typically the trigger algorithms are inclusive in ϕ, as the detector (and the known laws of physics) is symmetric in this direction.

The algorithms are designed to achieve maximal *trigger efficiency*, defined as the fraction of the events accepted by the trigger out of all events that the trigger path is targeting. Simultaneously, a trigger algorithm needs to have good enough *background rejection*, i.e. it needs to discard the majority of the events not targeted by it, in order to achieve a tolerable trigger rate.

Another desired property is a sharp *turn-on* of the efficiency, meaning that the efficiency rapidly rises from zero to its maximal value as a function of the true (transverse) momentum of the trigger object. Ideally, the trigger would always estimate the p_T of the object correctly, providing an efficiency curve that is a step function: objects with p_T below (above) the threshold are never selected (always selected). In reality, the limited momentum resolution can occasionally cause objects with too low (too high) momentum to be accepted (rejected). The sharper the turn-on is (with half-maximum close to the nominal threshold value), the closer we are to the ideal situation.

After the turn-on part, the efficiency curves typically reach a *plateau*, indicating that in the high-p_T range the trigger works at a stable maximum efficiency. In practice, the turn-on is often inspected as a function of the offline-reconstructed p_T by geometrically matching the trigger objects to the corresponding offline-reconstructed physics objects.

As the pileup increases, the probability of finding localized energy deposits in the detector that do not originate from the targeted hard interaction increases. Therefore for most trigger algorithms, discriminating between the targeted events (signal) and background events becomes more difficult as pileup increases, and higher trigger rates are needed to obtain the same efficiency as in lower-pileup conditions. If the trigger thresholds are kept constant, the rates typically grow at least linearly (sometimes even quadratically) with pileup. This pileup dependence can be partially mitigated by including pileup subtraction schemes in trigger algorithms. This aspect of algorithm development has become increasingly important in Run 2.

6.1.2 Trigger menu

The collection of trigger seeds which are running simultaneously during the data taking is called the *trigger menu*. A typical Run 2 L1T menu at the CMS contains 300–400 seeds. The maximum number of seeds allowed by the current architecture is 512, compared to 128 during Run 1. The total L1A rate, which is equal to the event readout rate, is determined by the rates from all the individual seeds. However, it cannot be calculated simply as a sum of the individual rates, because the same event can be accepted by several seeds.

As the readout rate is limited to 100 kHz, the trigger menu needs to be designed such that the total L1A rate stays below this value. This is done by carefully adjusting the thresholds used in each seed, such that the rate for each seed remains tolerable without compromising too much the sensitivity for the physics processes targeted by the seed. Additionally, seeds can be *prescaled*. If a seed is set to prescale N, part of the bunch crossings are masked from the trigger system, such that only every Nth event passing the thresholds defined for the seed is allowed to affect the L1A decision. Prescaling allows for example to collect (some) data with very low-threshold seeds that would produce a huge rate without prescaling.

In terms of trigger rate, about 3/4 of a typical CMS trigger menu consists of simple single- and double-object seeds, which require one or two high-p_T objects of the same type. Despite their simplicity, these seeds can trigger on a variety of interesting events

where heavy particles are produced. For example, a seed that requires two muons with high p_T can catch events where a Z boson decays into two muons, or a Higgs boson decays into Z Z, further decaying into two muons and two other leptons. A seed that requires large missing transverse momentum can catch events with high-p_T neutrino(s), such as the possible $H^{\pm} \rightarrow \tau^{\pm} \nu_{\tau}$ signal searched for in this thesis.

The remaining 1/4 of the total L1A rate is reserved for *cross seeds* that look for two trigger objects of different types, and for more complicated multiple-object algorithms that target very specific signals (such as Higgs boson production in vector boson fusion).

About half of the seeds are *unprescaled*, i.e. they have prescale equal to 1. They are used by physics analyses that benefit from minimizing the statistical uncertainty in the result by maximizing the amount of events they use. Prescaled seeds are used for physics analyses where high efficiency with low thresholds is more important than maximal amount of data. The collection of these unprescaled and prescaled "physics seeds" determines the limits of the physics reach of the experiment. In addition to the physics seeds, the menu can also contain additional seeds for calibration of the detector (e.g. energy scale calibration) and monitoring of the trigger system (including efficiency measurements). These seeds have usually low p_T thresholds, so they are prescaled.

The CMS trigger menu is adapted according to the changes in the LHC beam conditions. For example, sometimes the LHC delivers fills with a lower instantaneous luminosity than usual, and a menu with lower-threshold seeds or with smaller prescales can be used while still keeping the L1A rate below 100 kHz. Also during each fill, the instantaneous luminosity slowly decreases as the protons are consumed in the collisions. As the rate of hard collisions decreases with the luminosity, also the trigger rates go down. To benefit maximally from each fill, different versions of the menu called *prescale columns* can be activated on the fly while collecting data. When the L1A rate drops below a predefined threshold, a new prescale column with smaller prescale values for the prescaled seeds is activated, such that the rate returns close to 100 kHz.

Prior to running, the L1T rates can be estimated using samples of so-called *zero-bias* data. The zero-bias data is collected with an "agnostic" trigger that fires on any bunch crossing where a proton-proton collision is present, according to the structure of the LHC fill. In typical filling schemes, this corresponds to $\approx$30 MHz rate, so the zero-bias trigger always needs to be heavily prescaled while included in the menu.
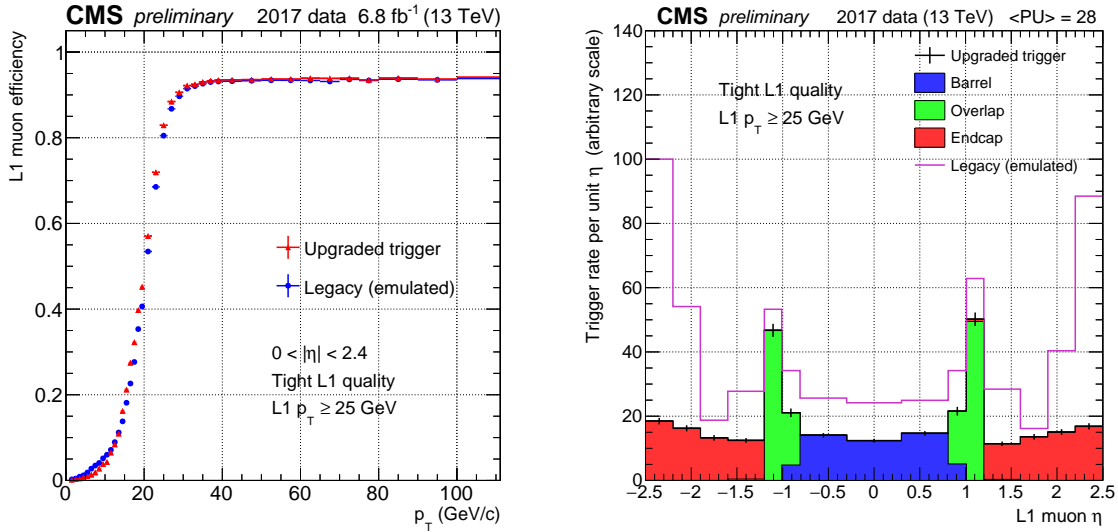
Figure 6.3: Single-muon efficiency of the Level-1 trigger system after the full Phase-1 upgrade (red), compared to the emulated performance of the old system used in Run 1 (blue). The upgraded muon trigger system provides a sharper turn-on and improved overall efficiency (left). Simultaneously, the trigger the rates are reduced in all track finder regions, as demonstrated by comparing the measured rates to the emulation of the Run-1 "legacy" trigger system (right). [121]

6.1.3 Muon trigger

As shown in Figure 6.2, the muon trigger takes input from all three types of muon stations: DTs, RPCs and CSCs. During Run 1, the muon tracks were reconstructed separately based on information from CSC, DT and RPC systems, and then combined in a global trigger module [119]. In the upgraded system, the information from different systems is combined already for the muon track reconstruction. The muon trigger system contains three separate *muon track finders*, corresponding to different $|\eta|$ regions in the detector. The *Barrel muon track finder* [122] takes input from DT and RPC chambers in the barrel region, i.e. $|\eta| < 0.83$ (as can be seen in Figure 5.7 in Section 5.5). Similarly, the *Endcap muon track finder* [123] uses the CSC and RPC chambers in the endcap regions, covering the $|\eta|$ range from 1.2 up to 2.4. They are complemented by the *Overlap muon track finder* [124], which takes input from all muon systems (DT, RPC and CSC chambers) in the overlap region, covering the $|\eta|$ region between the barrel and endcap muon track finder regions.

The barrel and overlap muon track finders are divided into 12 and 3 sectors in ϕ, respectively, where the muon triggering algorithms run in parallel. Each of the three track finders uses a different reconstruction algorithm, but they all aim to identify the

genuine muon tracks based on the measured hits in muon stations, and infer the muon p_T based on the track curvature, allowing the trigger to distinguish between high-p_T muons from interesting hard processes and a large background of low-p_T muons. The track identification is based on look-up tables containing predefined hit patterns to which the observed hits are compared.

The trigger primitives from the DT and RPC, namely the DT track segments and RPC hits, are merged into *superprimitives* before they are sent to the barrel muon track finder [122]. There the track reconstruction proceeds by extrapolation from inner muon stations towards outer stations, within acceptance windows defined by the superprimitives. The track p_T is determined based on correlations of the ϕ coordinates measured in different stations using a look-up table.

The unmerged DT trigger primitives are also sent to the overlap muon track finder, where they are combined with the RPC and CSC primitives. The overlap muon track finder algorithm combines reference hits in the inner layers with associated hits in outer layers, forming different track hypotheses which are tested against predefined *golden patterns* derived from simulation.

For the endcap muon track finder, the RPC hits are clustered by separate preprocessor boards, which combine the hits in adjacent strips into segments and assign ϕ and θ coordinates to them. These and the CSC segments are then used as input to the track finder, where the track reconstruction look-up table is based on boosted decision tree (BDT) regression. The BDT is trained with simulated single-muon events and uses 11 input variables from the RPC and CSC segments to map them to muon tracks.

For each event, each of the three track finders sends (at maximum) 36 muons to the *global muon trigger*, where the information is combined, and duplicates are resolved. The global muon trigger then sends eight highest-p_T and highest-quality muons to the global trigger.

The performance of the upgraded muon trigger system is illustrated in Figure 6.3. Compared to the emulated performance of the Run-1 muon trigger, the upgraded one provides a sharper turn-on with a good overall efficiency, while also considerably reducing the trigger rates across all track finder regions.
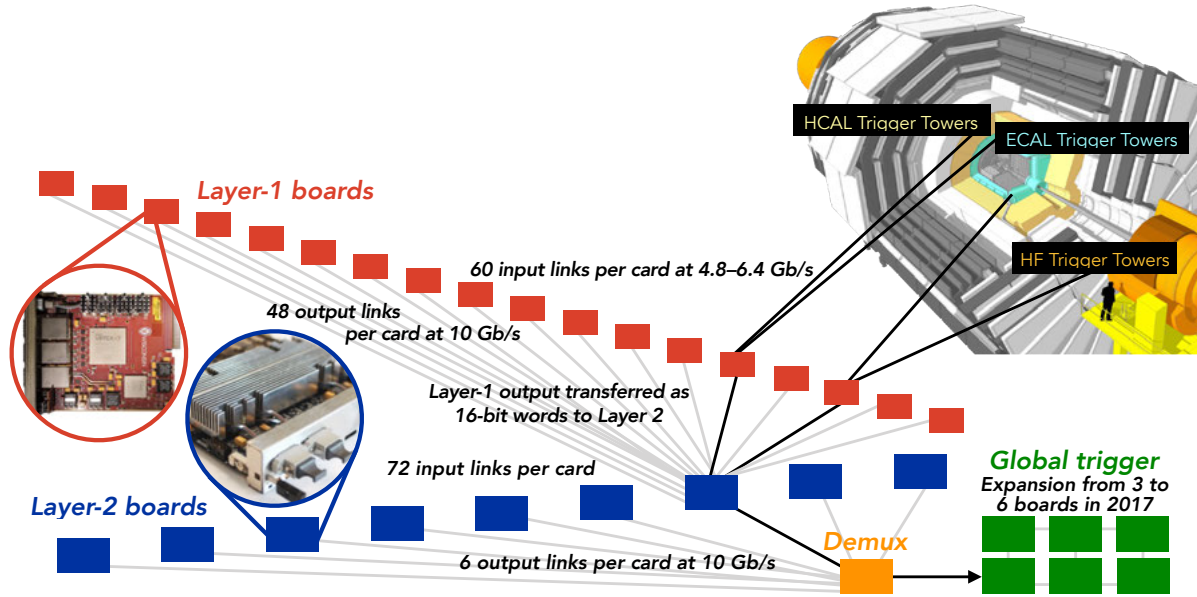
Figure 6.4: The architecture of the upgraded calorimeter trigger system. In the first layer, 18 boards pre-process the trigger tower information, sent to the second layer via fast optical links. The second layer functions in a time-multiplexed way, so each full event is processed in one of the nine boards. An additional "demux" card forwards the information from the second layer to the global trigger.

6.1.4 Calorimeter trigger

Both ECAL and HCAL provide input to the calorimeter trigger. They have dedicated front-end electronics to shape the signal pulses, digitize them and sum them together geometrically into trigger primitives, referred to as *trigger towers* (TT). In the barrel, the trigger towers have $\eta \times \phi$ granularity of approximately 0.087×0.087, corresponding to 5×5 ECAL crystals and one HCAL tower behind them. In the endcaps, the trigger towers are larger, with sizes of up to 0.17×0.17.

The two-layer architecture of the upgraded calorimeter trigger is illustrated in Figure 6.4. The first layer consists of 18 processor boards, which receive the TTs from a given geometrical region. They take care of TT-level pre-processing of the data, such as energy calibration of each tower, summing the transverse energies (E_T) of the TTs in the given region, and calculating the H/E ratio, i.e. the ratio of HCAL and ECAL energy deposits. After the first layer, the data corresponding to one event are collected from all boards and sent to one of the boards in the second layer. To transform information efficiently, the TT information is formatted into compact 16-bit words in the first layer, and transferred to the second via high-speed optical links.

The second layer consists of nine processor boards, and works in a *time-multiplexed* manner. In a time-multiplexed trigger, each board has access to a complete event, and can evaluate it using information from the full detector, i.e. a full set of TTs and their properties calculated in the first layer. As transferring the full detector information to the second layer takes typically ≈ 7 bunch crossings, nine cards are enough to smoothly process all the incoming data. The time-multiplexed approach is used in CMS also in DAQ and HLT systems, but the upgraded calorimeter trigger system is its first application in L1T.

The boards in the second layer host the algorithms that reconstruct and identify particle candidates and calculate the global energy sums, such as H_T, p_T^{miss} and H_T^{miss}. For each event, the second layer identifies up to 12 jets, electron and photon (e/γ) candidates and τ_h candidates, and determine their p_T as well as η and ϕ coordinates. The sophisticated reconstruction and identification algorithms developed for this task are detailed in Ref. [125], and in the following their main features are presented.

The jet reconstruction is seeded by an energy deposit in a TT that exceeds a pre-programmable threshold of a few GeV. The jet energy is then determined by summing the TT energies in a 9×9 TT window centered around the local maximum. An inequality mask is applied to ensure that the seed is a local maximum and to prevent jet double-counting. The window size of 9×9 TTs corresponds to the cone radius of 0.4 used in the offline jet reconstruction (Section 8.2.3). The effect of pileup is estimated with a "chunky donut" algorithm [125], which sums the TT energies in four 3×9 strips next to each edge of the 9×9 TT jet. The TT energies of the three lowest-energy strips (9 TTs in total) are summed and subtracted from the jet energy. Finally, the jets are calibrated as a function of the p_T and η using a dedicated look-up table. H_T is calculated as the jet p_T scalar sum, and H_T^{miss} as the negative jet $\vec{p}_T$ vector sum.

As the electrons and photons cannot be distinguished without the tracker information, they are treated commonly as e/γ candidates. They are identified using a *dynamic clustering*: a cluster is seeded by a TT with $E_T > 2\,\text{GeV}$, and built by iteratively adding all new neighboring TTs with $E_T > 1\,\text{GeV}$ to the cluster. The clusters are allowed to contain at most 8 TTs, and they are allowed to have a wider spread in ϕ (due to bending of trajectories in the magnetic field) than in η. Finally, the clusters are categorized based on the distribution of energy in them, and only those compatible with the e/γ hypothesis are accepted.
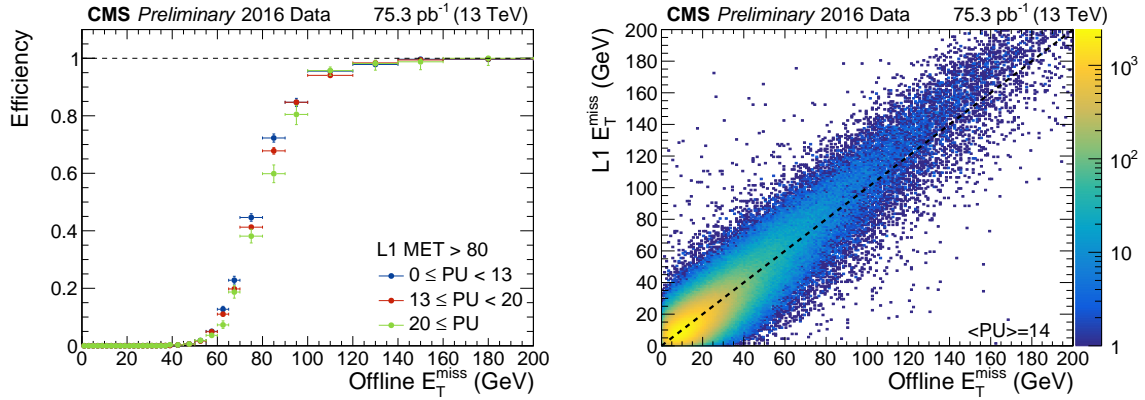
Figure 6.5: The turn-on curve of the trigger efficiency for the Level-1 trigger path requiring the missing transverse momentum of $> 80\,\mathrm{GeV}$, as a function of the offline-reconstructed $p_\mathrm{T}^{\mathrm{miss}}$ (left), and the correlation between the Level-1 and offline-reconstructed missing transverse momentum (right). [127]

The τ_h candidates are identified using a dedicated clustering algorithm, similar to the e/γ algorithm. As many final states of the hadronic τ decays contain neutral pions (Section 2.5.1) which in turn decay to photons, the fingerprint of a τ_h decay in the calorimeters can correspond to several e/γ clusters. For τ decays with three charged hadrons, the energy deposits are even more spread out in ϕ. Hence the τ_h identification algorithm is allowed to merge together several energy clusters if certain proximity conditions are fulfilled [126]. For both e/γ and τ_h candidates, the position is determined as an energy-weighted average of the TT positions, yielding a factor of four improvement in position resolution compared to Run-1 algorithms that used the position of the seed TT.

The e/γ and τ_h candidates can be required to be isolated. The isolation energy sum is calculated in a 6×9 window centered around the candidate, and the reconstructed candidate energy is subtracted from the sum. For optimized performance, the energy thresholds for passing the isolation are stored in a look-up table, and they depend on the candidate p_T, η, and the number of non-zero trigger towers in the event, which reflects the amount of pileup. The isolation provides efficient discrimination between τ_h candidates and quark and gluon jets misidentified as τ_h candidates. Both e/γ and τ_h candidates are calibrated using dedicated look-up tables, allowing energy corrections that depend on the reconstructed p_T, seed tower η and cluster shape variables.

The Level-1 $\vec{p}_\mathrm{T}^{\mathrm{miss}}$ is reconstructed as a vector sum of all TT energies that exceed predefined thresholds. For optimized pileup mitigation, the thresholds depend on

η and the actual size of a TT in terms of $\Delta\eta \times \Delta\phi$. As the analysis presented in this thesis targets a final state with large missing transverse momentum, the Level-1 trigger algorithm used for this analysis requires that $p_T^{miss} > 80\,\text{GeV}$. The turn-on of this trigger algorithm a function of the offline-reconstructed p_T^{miss}, as well as the correlation between the Level-1 p_T^{miss} and offline-reconstructed p_T^{miss}, are shown in Figure 6.5.

The last piece of the calorimeter trigger is the de-multiplexing ("demux") board, which collects the data from the second layer, merges the partial vector sums calculated in the second layer, and forwards the data to the global trigger. A more detailed description of the updated calorimeter trigger and its hardware is given in Refs. [125, 128–130].

6.1.5 Global trigger

The global trigger collects and combines the information from the global muon trigger and from the demux board of the calorimeter trigger. Upon arrival to the global trigger, these input data are synchronized to each other and to the LHC clock. The global trigger contains the active trigger menu (including the chosen prescales), and uses it to evaluate each event, and to perform the final decision to accept or discard it.

The global trigger takes into account additional criteria known as *trigger rules* before sending the L1A signal. Trigger rules are pre-programmed settings that are designed to prevent buffer overflows in the DAQ system that would cause serious dead time. For example, the trigger rules only allow one L1A signal per three subsequent bunch crossings, so that each accepted bunch crossing is always followed by two untriggered ones, ensuring at least 75 ns spacing between two accepted events. The trigger rules create some dead time themselves, but it is typically $< 1\%$. The global trigger also takes input from several beam monitoring systems, in order to only accept bunch crossings with a genuine proton bunch crossing.

Originally, the global trigger operated on a single board. In 2016 it was extended to three boards, and in 2017 to six boards. The larger amount of logic resources in the global trigger allows implementation of complex trigger paths that have traditionally been implemented only at HLT level. For example, invariant or transverse masses can be calculated for combinations of different objects, and complicated cross-seeds that require a specific combination of different types of trigger objects can be implemented. This also allows introduction of trigger paths that target specific event topologies. For example, a specific trigger algorithm that targets Higgs boson production via vector boson fusion based on the invariant mass of jet pairs has been developed [131].

Figure 6.6: Examples of the data quality monitoring plots, used both for online monitoring of
the trigger system and express certification: a normalized p_T distribution of L1T
muon objects as obtained from the global muon trigger (left), and (η, ϕ) occupancy
of isolated e/γ candidates as obtained from the second layer of the calorimeter
trigger (right).

The global trigger can also be used to produce in-situ estimates of the trigger rates pro-
duced by different prescale columns, which facilitates the optimal usage of prescales
under changing LHC beam conditions.

6.1.6 Operation and monitoring

During data-taking, the L1T system is continuously monitored by the *trigger shifter*,
a dedicated member of the shift crew that operates the detector. The trigger shifter
adjusts the prescale columns based on the instantaneous luminosity, and oversees the
total L1A rates as well as the individual rates from the different seeds in the menu.
In cooperation with the *data quality monitoring* (DQM) shifter, the trigger shifter also
monitors a collection of continuously updated DQM plots. The aim is to identify any
problems that might prevent successful data-taking or compromise the data quality,
and take immediate action to solve them.

The DQM plots used to monitor the L1T show e.g. the trigger primitive occupancies in
calorimeters and muon stations, kinematic distributions (p_T, η, ϕ) for different types
of trigger objects, and timing information for selected trigger seeds. Monitoring of the
timing is critical, since if the L1A signal is sent too early or too late, the data is read
out at the wrong bunch crossing, and the targeted event is lost. Examples of the L1T
DQM plots are shown in Figure 6.6.

The trigger decisions made in the L1T hardware are also constantly emulated with a bit-level emulator, which is also used to estimate the effect of the trigger in the simulated events. During data-taking, the agreement between the emulated and observed distributions for different trigger objects is monitored by the DQM shifter. This cross-check ensures that the trigger system behaves as expected, and on the other hand it confirms that the emulation code models the real system in a correct way.

The shift crew, including the trigger shifter, is supported by a group of *on-call experts* for different subsystems. Abnormal trigger rates or distributions can indicate a problem in the trigger system. In such cases, the shifters promptly contact the on-call experts, who proceed to diagnose the system (usually remotely), and decide on actions to be taken. The L1T on-call expert is further supported by a number of *trigger object experts* who are typically experienced scientists who have contributed to the development of the trigger system and the algorithms running on it. Both the shift crew and the on-call experts document their observations and actions in electronic logbooks. The timestamps of logbook entries can be used to match them to the collected data.

6.1.7 Data certification

After the data has been successfully recorded, it is inspected more carefully in the *offline certification* process. During Run 2, a new two-step process was established for data certification, with significant contributions by the author who served as the CMS Level-1 trigger offline data certification co-coordinator in 2018. The first step, the *express certification*, is performed almost immediately after the data taking. It is followed by a more thorough *final certification* once the recorded events have been fully reconstructed. The certification is performed for all *runs* (continuous sequences of data taking) recorded while the LHC is delivering collisions, and during which all subdetectors are powered on and set to acquire data (since input from all subsystems is required for successful offline reconstruction of the full event).

The express certification serves as a first check and is carried out for the collected data almost immediately after data-taking (typically within 24 hours), in order to identify any abnormal behavior in the trigger system that might have passed unnoticed despite the online monitoring efforts. The goal is to identify and fix any problems before they affect long periods of data taking.

In the express certification, the behavior of the total L1A rate is examined as a function of time. Changes in rate are compared to information about the beam conditions,

Figure 6.7: The time evolution of the total L1T output rate during a typical proton-proton colli-
sion run. The rate decreases gradually as the instantaneous luminosity decreases.
The sharp rises around luminosity section (LS) 200 and 550 correspond to changes
of prescale column. The rate drop around LS 900 might indicate a trigger problem,
and would be inspected as part of offline certification. In this particular case, it
corresponds to a drop in instantaneous luminosity.

prescale columns applied, status of each subdetector, and dead time. The aim is to
understand all changes in the trigger rates, since unexpected changes in rate can
indicate a problem in the L1T itself or in one of the subdetectors sending input to
it. For example, a malfunctioning detector element can either cease to send trigger
primitives, causing a drop in trigger rates, or produce a lot of noise interpreted as
trigger primitives, causing a peak in the rates. Logbook entries written down by the
shift crew and on-call experts during data can be used to provide more information
when needed. An example of the time evolution of the total L1A rate during a proton-
proton collision run is shown in Figure 6.7.

In addition to the total L1A rate, the rates of selected individual trigger seeds that
target different physics objects are checked. The individual rates are compared to
reference values obtained from previous successful runs. Since the rates depend on
pileup, which can vary during the data taking considerably (mostly depending on
the instantaneous luminosity), the reference rates are specified as a function of pileup.
Also in the case of unexpected changes in the total L1A rate, tracking down its origin
to individual trigger seeds often helps in understanding the origin of the changes.

For each run, the DQM plots monitored by the shift crew already during data taking
are also double-checked as part of the express certification, including the comparison
of emulated and recorded trigger objects.

The *final certification* is performed after the recorded data has been fully reconstructed for the first time. This *prompt reconstruction* typically takes place within 48 hours from data taking. The primary goal of the final certification is to select the high-quality data suitable for physics analyses, and discard the data that do not meet the quality criteria. A second goal to discover any problems that might have passed unnoticed both during data taking and in express certification, so that they can be understood and prevented in the future.

In the final certification, the efficiencies for different types of trigger objects, calculated using tag-and-probe method [132], are inspected as a function of p_T, η and ϕ. Also the resolutions are studied by comparing these variables between trigger objects and geometrically matched offline-reconstructed physics objects. If the obtained efficiencies and resolutions match to the expected performance, and the results of the express certification suggest that the trigger system operated correctly, the data is certified as valid to be used in physics analyses (from the L1T point of view).

The data collected during a CMS run are divided into smaller units known as *luminosity sections*. One luminosity section corresponds to 23.3 seconds of data taking. To maximize the data available for physics analyses, the final certification is performed per luminosity section. If the data from some run do not meet the quality criteria, the run is inspected at luminosity section level to identify the ones affected by the performance loss, and only the data attributed to those luminosity sections are discarded.

Both in the case of express and final certification, when a certain run does not pass the certification criteria, the certification experts and the detector operation experts work together to identify and understand the performance loss, and take any necessary action to prevent it from happening again.

In 2016–2018, the overall CMS data certification efficiency, i.e. the fraction of collision data certified as good for physics analyses in terms of integrated luminosity, has been ≈95% each year. The remaining 5% of data has been discarded because one or several subsystems have decided to discard them based on their certification processes. Of the 5% of data discarded, the L1T certification process has been involved in 0.5–1.5 percent units, depending on the year. The majority of these data discarded by the L1T has also been discarded by at least one other subsystem, typically indicating that a large detector malfunction affecting several subsystems occurred. The data losses exclusively due to L1T-related issues is at per mille level. In conclusion, the upgraded L1T system operated in a robust and reliable way throughout Run 2.

6.2 High-level trigger

The high level trigger is a fully software-based trigger system, running on the *event filter farm* above the CMS detector, built of commercial computers, which are mounted on racks. In total, it contains approximately 30 thousand processor cores. As shown in Figure 6.1, the HLT takes as input full event data as read out from all CMS subdetectors (at a maximum rate of 100 kHz), performs a fast reconstruction of the events and makes a trigger decision for each event in a few hundred milliseconds, producing an output rate of a few hundred Hz.

The HLT output rate is constrained by the bandwidth available for data transfer to "Tier 0" storage elements, limited to ≈ 5 Gb/s, corresponding to a maximum event rate of a few kHz. A more stringent constraint arises from offline computing resources: normally a first round of offline reconstruction is performed for all recorded events within 48 hours from data-taking. This *prompt reconstruction* makes it possible to inspect the quality of data in its final form, in order to ensure that the detector works correctly. The computing resources available for the prompt offline reconstruction limit the HLT output rate to < 1 kHz.

Workaround methods have been developed to avoid both of these constraints in specific cases where additional output rate is needed. In *data scouting*, only the HLT-level variables are stored instead of full event information (avoiding the first constraint). On the other hand, *data parking* means that data are saved without performing the prompt reconstruction (avoiding the second constraint), and the events are reconstructed later during technical stops when the computing resources are available.

Although the architecture of the HLT is completely different from the L1T, the basic concepts apply to both. Also the HLT contains a trigger menu with hundreds of different seeds, which are often called *trigger paths* in the context of the HLT. The kinematic thresholds defined for each path are limited by the total HLT output rate, but also by the kinematic thresholds used in Level-1 trigger.

Unlike the Level-1 trigger, the HLT benefits also from the tracker information. The tracks are reconstructed following the iterative tracking approach described in Section 8.1.1. In the HLT version of the algorithm, the last and most resource-consuming iteration steps are limited to certain regions of interest, indicated by the presence of calorimeter jets or other tracks.
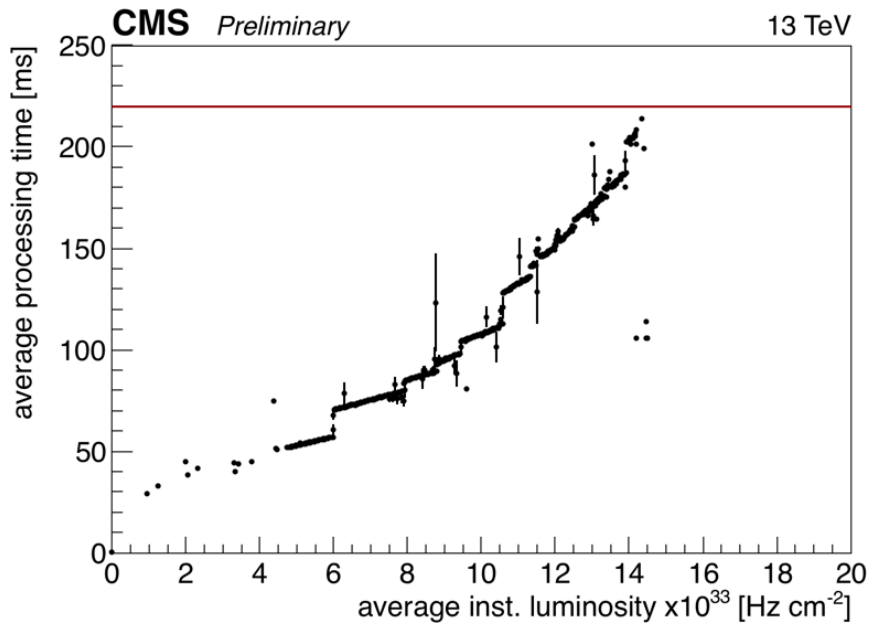
Figure 6.8: Average HLT processing time per event as a function of instantaneous luminosity, measured during a specific LHC fill (5393) in 2016. The red line corresponds to the maximum allowed processing time in 2016. [133]

To benefit from the full event information, the HLT runs a streamlined version of the Particle Flow (PF) event reconstruction algorithm that is also used in the offline event reconstruction and presented in detail in Section 8.1. The HLT version of the algorithm has been modified to meet the stringent limit of $\approx$200 ms of processing time per event. The data processing steps are divided into independent reconstruction and filtering modules, following each other such that the majority of events can be filtered out as early as possible. The simplest and fastest modules, such as processing of calorimeter information, are executed first, and as soon as an event fails one filtering step, its processing is halted completely. This way the relatively slow full PF reconstruction is actually performed only for a tiny fraction for events.

During Run 2, the HLT hosts $\approx$400 different trigger paths, most of which rely on HLT-level PF candidates: muon, electrons, photons, neutral and charged hadrons. The PF candidates are clustered into jets and used to construct other higher-level HLT objects such as τ_h candidates. Jets originating from bottom quarks (b jets) can be identified at the HLT with the combined secondary vertex algorithm, the offline version of which is presented in Section 8.2.4. Usage of the tracker information enables discrimination between electrons and photons in the HLT, allowing for example specific trigger paths that target Higgs boson decays to photons.

The HLT paths typically request the presence of one or several HLT objects of specific types, above certain p_T thresholds and within the desired $|\eta|$ range. The menu includes also some more complex paths, comparable with analysis-specific offline selections. The usage of the PF algorithm also means that the advanced pileup mitigation techniques developed as part of the offline PF algorithm can be used.

A typical example of the HLT performance in 2016 is shown in Figure 6.8, where the average processing time in the HLT per event is shown as a function of instantaneous luminosity. The red line corresponding to 220 ms processing time indicates the maximum time limit allowed for the HLT in 2016.

Chapter 7

Simulation of collision events

Computer simulation of collision events and the detector response is an indispensable part of the methodology in modern particle physics. The methods developed over the past thirty years allow us to perform very detailed and accurate, but often computationally expensive, simulations. The same basic methods are used to cover several orders of magnitude in energy and momentum, from keV scale to TeV scale.

Results from simulations are used in several stages of experimental particle physics research projects. Firstly, they are used in designing new experiments or upgrading existing ones to study if the desired sensitivity and precision can be achieved with a planned architecture. Secondly, they are used for feasibility studies, e.g. to check if a specific decay channel of a particle is worth a full analysis effort. Thirdly, simulated events are commonly used to derive signal hypotheses, such as the charged Higgs boson signal searched for in this thesis. Similarly, simulations provide a way to estimate background contributions from *irreducible background processes* that are difficult to measure from data without signal contamination. Development of algorithms used to identify physics objects, such as those presented in Section 8, also relies on simulated events. These include modern multivariate analysis classifiers, where simulated event samples are used to train the classifiers. Additionally, comparisons of kinematic distributions between data and simulation are often useful in the validation of analysis methods. Finally, simulations can provide estimates needed to interpret the analysis results, for example to calculate a full cross section from a measured fiducial cross section with limited phase space, or to infer particle composition of jets from the measured jet properties.

In this chapter, the key concepts of simulation of proton-proton collision events at the LHC are presented. Emphasis is given to tools and methods used to produce the

signal and background estimates used in the $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ analysis presented in this thesis. These simulated events are stored in the same format as the real data events and treated in the analysis identically to the data.

As discussed in Chapter 4, the difference between the cross section of a specific process of interest, such as a Higgs boson production process, and the total inelastic proton-proton cross section is typically many orders of magnitude–even a factor of 10^{15}. Therefore it is computationally impossible to obtain large samples of specific types of events by simulating what the LHC does in real life, i.e. by simply generating huge amounts of generic proton-proton interaction events, very rarely producing an event of interest. Instead, the event generation typically starts with the description of the *hard interaction* happening via a specific process of interest. Then the outcoming particles are evolved forward in time step by step, while the incoming particles are evolved backward, towards the initial state that preceded the interaction.

The workflow used to produce simulated pp collision events can be factorized into four separate steps: the *generation of the hard process* (using perturbative QCD) is followed by the forward and backward evolution of *parton showers*, the *hadronization* of the partons produced, and modeling of *secondary interactions* of the partons in the protons. In Section 7.1, these basic steps of *event generation* are outlined, and some commonly used algorithms and software packages are presented. A more comprehensive description of all these steps can be found in Ref. [134]. To model the response of the experimental apparatus, such as the CMS detector, the generated events are processed through *detector simulation*. In Section 7.2, the simulation of the CMS detector is discussed. The subsequent steps of pileup simulation and signal digitization are covered in Sections 7.3–7.5.

7.1 Event generation

The four steps of event generation—the hard process, parton showers, hadronization and secondary parton scatterings—are illustrated in Figure 7.1. A variety of software exists to carry out the different steps. The *general-purpose event generators*, such as HERWIG [135] and PYTHIA [136], are equipped with algorithms to perform all the steps. In some cases the modeling of the hard process requires separate software, such as MADGRAPH [137] or POWHEG [138–140], which are commonly used to produce parton-level events. Their output is then directed to a general-purpose generator software, such as HERWIG or PYTHIA, to perform the subsequent steps. For example,

Figure 7.1: Sketch showing the different steps of simulation of a proton-proton collision with a Monte Carlo event generator. The red circle represents the hard interaction of two partons (blue lines), surrounded by red parton showers. The partons become confined into hadrons (light green blobs) that further decay into lighter particles (dark green blobs). The purple blob indicates a secondary parton interaction that also yields parton showers that eventually hadronize. [142]

Les Houches Event (LHE) file format [141] can be used to describe parton-level events, providing them as input to a general-purpose software for parton showering and hadronization.

7.1.1 Hard process

For hard processes with large momentum transfer, the asymptotic freedom of QCD allows perturbative treatment and the hard interaction can be conveniently described with a Feynman-diagrammatic approach. The *Feynman amplitude* of the process can be calculated by summing the terms corresponding to the Feynman diagrams that describe the process up to a chosen order in perturbative QCD (traditionally at leading order, these days increasingly commonly also at next-to-leading order). The square of the amplitude gives the *matrix element* for the process, which can be used to calculate differential cross sections that determine the distributions of kinematic variables.

Integrating over a differential cross section, the total cross section for the process is obtained, which determines the correct normalization of the simulated sample. The cross sections depend on the center-of-mass energy of the initial particles present in the collision.

The general-purpose generators such as PYTHIA contain a wide selection of LO matrix elements for different $2 \to 1$, $2 \to 2$ and $2 \to 3$ processes in the SM and in selected extensions of it. MADGRAPH can be used to generate events at LO, or at NLO using its recent extension MADGRAPH5_aMC@NLO [137], in a generic way for any process that can be described using the Universal FeynRules Output (UFO) format [143]. In many cases, the LO generation can provide correct predictions for shapes of kinematic distributions, but the overall normalization of the process (the inclusive cross section) can be clearly off due to lack of higher-order corrections. In such cases, the normalization is corrected by a multiplicative factor calculated from theory, referred to as a *K-factor*.

The evaluation of the cross sections involve estimation of high-dimensional phase space integrals. For this reason, Monte Carlo techniques, such as the *hit-and-miss* integration method, are used in event simulation to such an extent that the resulting simulated events are often referred to as Monte Carlo events. If an event has n particles in the final state, the phase space has $d = 3n - 4$ dimensions (components of three-momenta of the n outgoing particles minus the constraints from four-momentum conservation). Monte Carlo integration converges as $1/\sqrt{N}$ where N is the number of iterations, independent of the dimensionality d, whereas other numerical integration methods converge as $1/N^{C/d}$. The constant C depends on the method, but in all non-Monte-Carlo methods larger dimensionality implies slower convergence. While integrating over the phase space, cuts are often introduced for phase space parameters to take into account the detector acceptance (in terms of energy thresholds or geometrical limits of a detector), or to avoid unphysical divergences in the cross section. The latter requires introduction of renormalization and factorization scales, which are usually set to the same value.

7.1.2 Parton distribution functions

In the proton-proton collisions, the center-of-mass energy of the partons in the hard collision varies from event to event according to *parton distribution functions* (PDFs). The PDF $f(x, \mu)$ describes the probability that a certain type of quark or gluon is found carrying a fraction x of the total momentum of the incoming proton when

Figure 7.2: Examples of NNPDF3.0 parton distribution functions, shown as a function of x at low momentum transfer of $10\,\text{GeV}^2$ (left) and at high momentum transfer of $10^4\,\text{GeV}^2$ (right), with $\alpha_s(M_Z^2)$ set to 0.118. [149]

it is probed at an energy scale μ. While the scale dependence of the PDFs can be calculated theoretically with so-called DGLAP equations [144–147], their functional form is determined by the non-perturbative physics inside the proton, so it needs to be fitted from data.

The choice of PDF affects not only the total cross section for the process, but also the kinematic distributions of final-state objects and the global event shapes. The PDFs enter not only the hard process, but also in the simulation of parton showers and secondary parton interactions.

During the LHC Run 2, a recent set of PDFs produced by the NNPDF Collaboration, known as NNPDF3.0 [148], is widely used in event generation. These PDFs are based on a large variety of experimental results from different experiments operating at different energy scales, and they are scrutinized with several closure tests. Examples of the NNPDF3.0 parton densities are shown in Figure 7.2.

7.1.3 Parton showering

Analogously to how charged particles can emit bremsstrahlung photons in QED, colored partons can radiate gluons in QCD. But due to the non-Abelian nature of the QCD, gluons can also themselves radiate gluons (while photons cannot radiate other photons). These "splittings" give rise to a cascade of partons where they are continuously scattered, annihilated, and new ones are produced.

To model these *parton showers*, results from perturbative QCD are applied in approximative ways to implement algorithms that create a shower as a series of probabilistic time-ordered $1 \rightarrow 2$ splittings. The process can be seen as a stepwise Markov chain, so it can be efficiently implemented in simulation using Monte Carlo techniques.

In the parton showering step, the initial-state partons are showered by *backward evolution* towards the incoming protons, describing the evolution of the PDFs as a function of energy scale in a time-reversed way, from the high momentum transfer scale μ of the hard process to a lower scales characterizing the partons in the incoming protons prior to the collision [150]. The final-state partons are evolved *forward* in time, from the high-μ collision gradually down to an *infrared scale* of ≈ 1 GeV, where the non-perturbative confinement effects step in and the *hadronization* takes place. Going down in μ corresponds to an evolution from short to long distance scales. The key task in the parton shower modeling is the evaluation of *Sudakov form factors* that give probability for a parton to survive, i.e. not to undergo a splitting, between two scales.

The implementations of parton showering differ between generators. HERWIG uses angular-ordered parton showers, where the splittings with largest angles are performed first to ensure correct treatment of soft gluons. PYTHIA uses p_T-ordered dipole showers, where the hardest emissions come first. The two orderings are considered to be effectively equivalent. A more detailed description of the parton showering models and the related calculational schemes can be found in Ref. [142].

7.1.4 Matching and merging

The simulation of hard process and parton showers complement each other. In the hard process, high-p_T partons with large angular separation are modeled well, whereas soft and collinear partons are not—in the parton showers it is vice versa. From the point of view of perturbative QCD, parton showering introduces higher-order corrections to the hard process cross section calculated at a fixed order.

To a first approximation, the details of the implementation of the parton showering alter the shapes of the distributions sensitive to the parton showering part of the process, such as the jet multiplicity distribution, but do not affect the overall cross section, which determines the normalization of the simulated sample. However, as changes in the distributions affect the acceptance of the events with respect to online and offline selections, care must be taken to implement the parton showering such that the shape of these distributions is as correct as possible.

As the hard process cross sections are *inclusive* in multiplicity of final-state partons, while the parton showers produce *exclusive* final states with a specific multiplicity, the two approaches must be combined in a careful way to avoid double-counting or undercounting of some final states.

Several approaches exist to ensure a correct transition from the hard process to the parton shower. In *matching* methods, the higher-order corrections are provided by the parton shower, so the corresponding contribution is subtracted in the hard process calculation. In *merging*, the hard process is calculated exclusively at each final-state parton multiplicity, including partons above a given kinematic threshold in the matrix element, and the parton showers containing soft partons under that threshold are then mapped to them accordingly. In a p_T-ordered parton shower this means that the hardest emission, which mostly affects the structure of the final state, is handled at the level of the hard process. In practice, the two approaches can also be combined by merging matched results.

Both POWHEG and MADGRAPH5_aMC@NLO provide methods to correct the first (highest-p_T) emission in the parton shower at the level of hard process. In POWHEG, this is done by weighting the matrix element with a suitable NLO-to-LO correction factor, while in MADGRAPH5_aMC@NLO the matrix element is split into a sum of LO and NLO contributions. The latter approach yields event samples that contain simulated events with both positive and negative weights, meaning that larger samples need to be generated to achieve the desired statistical precision.

7.1.5 Hadronization

When the momentum transfers reach the $\approx 1\,\mathrm{GeV}$ scale, the perturbative QCD breaks down, and the further evolution of the event needs to be handled with a nonperturbative model.

Figure 7.3: Illustration of the string hadronization model. The potential between a quark (blue line) and an antiquark (green line) is represented as a stretching string (red line). In the picture, the vertical axis corresponds to time (t) and the horizontal axis to a spatial dimension (z). As the string gets split into several smaller strings over time, quark-antiquark pairs recognized as hadrons are created. [142]

Hadronization models describe how the final-state partons are confined into colorless hadrons, and how the initial-state partons interact with the rest of the protons in the incoming beam. The hadronization models are inspired by QCD, but they cannot be rigorously derived from theory due to its nonperturbative nature. In these models, partons are no more treated independently, but instead a color-connected parton system hadronizes collectively. Due to their effective nature, they contain free parameters that need to be tuned to produce output that matches experimental results.

HERWIG uses a *cluster hadronization model*, which relies on a property of QCD known as preconfinement [151]. In the cluster model, first color-singlet combinations of partons called clusters are formed, and then the clusters decay into two hadrons that can further decay into stable hadrons. On the other hand, PYTHIA applies a *string hadronization model*, also known as the Lund model, illustrated in Figure 7.3. It is based on the fact that in QCD, the potential between two quarks rises linearly with their distance, stretching a "color string" between them. Thus the hadronization is modeled in terms of massless relativistic strings, corresponding to quark-antiquark pairs, which can split into two when stretched enough by the two distancing quarks.

The hadronization process often produces short-lived resonances. For computational efficiency, the decays of heavy resonances are typically factorized into a separate step. For resonances with narrow widths, the decay can be simulated in the rest frame, followed by a Lorentz boost back to the laboratory frame.

7.1.6 Secondary parton interactions

As the total proton-proton cross section is dominated by soft QCD events, soft secondary interactions commonly occur between partons of the colliding protons, in addition to the hard interaction described in the previous steps. The soft pp interactions consist mostly of diffractive scattering and multiparticle production processes that produce low-p_T particles.

As the associated momentum transfers are low, again the perturbative QCD is not applicable and effective models are needed. Therefore nonperturbative *multiple-parton interaction models* and *diffraction models* with tunable parameters are used to describe the secondary interactions. A review of the currently used models is provided in Ref. [134]. The final-state particles that are not produced in the hard process are referred to as the *underlying event*. Eventually, these particles are also subject to hadronization described in Section 7.1.5.

7.2 Detector simulation

After all steps of event generation have been carried out, *radiation transportation* software is used to simulate the propagation of particles through the detector material. In the LHC experiments, including CMS, the GEANT4 toolkit [152, 153] is used to model the interactions between generated particles and the detector.

A detailed description of the geometry and materials of all components of CMS is implemented in GEANT4, including both the active detector elements and the passive material such as cables and cooling systems. The software traces the particles through the detector in small steps, using Monte Carlo simulation to impose the particles to different stochastic processes according to their probabilities.

GEANT4 includes models to describe a variety of interactions with the detector material, including the effects of electric and magnetic fields, bremsstrahlung, photon conversions, multiple scattering, ionization, and interactions between hadrons and nuclei ranging from MeV-scale elastic scattering of neutrons to GeV or even TeV scale hadron showers. Most interactions are not modeled strictly from first principles. Instead, approximative semiempirical schemes are used, with free parameters that are tuned to calibrate the models to produce realistic results.

7.3 Pileup modeling

The effect of pileup is added in the simulated collision events by mixing them with separately simulated *minimum-bias* events, i.e. soft QCD interactions that dominate the total inelastic proton-proton interaction cross section. For this purpose, a sample of minimum-bias events is generated using PYTHIA and processed through all simulation steps as described in Sections 7.1 and 7.2 up to the radiation transportation modeling with GEANT4. After this step, the simulated energy deposits (hits) from the minimum-bias events are mixed with the ones from the original hard collision event.

The amount of minimum-bias events to be mixed with the original event is determined as follows. First, a pileup distribution, i.e. the distribution of the *mean number of interactions per bunch crossing* ($\langle\langle N_{PU}\rangle\rangle$), is specified. The $\langle N_{PU}\rangle$ distribution is based on the distribution observed in real data, if available, or on an expected distribution if simulation is performed before collecting data. Then, for each generated hard event, the mean number of interactions per bunch crossing is sampled from the $\langle N_{PU}\rangle$ distribution. The amount of pileup in the specific event at hand is sampled from a Poisson distribution with a mean set to the chosen value of $\langle N_{PU}\rangle$.

The $\langle N_{PU}\rangle$ distribution includes both in-time and out-of-time pileup. To model the out-of-time pileup correctly, several bunch crossings around the nominal one are included in the simulation by varying the time information associated with the minimum-bias events accordingly. Thus typically hundreds of minimum-bias events are needed to model pileup for one generated hard collision event.

7.4 Digitization

After the GEANT4 simulation has described the local energy deposits ("hits") in the active detector elements, they are *digitized*, i.e. converted to electric signals. The digitization is performed with a separate module in the CMS software, emulating the behavior of the readout electronics that are used in the real data acquisition. The parameters of this conversion are tuned to obtain as realistic detector response as possible, including the electronics noise.

At this stage, also the emulation of the Level-1 trigger electronics is performed. This way, the simulated events can be stored in a format that is equivalent to the one used for data as they are read out from the detector. Every simulated event is also equipped

with the information about which hardware trigger criteria (Level-1 trigger seeds) it would have passed if it were a real event.

7.5 Fast simulation

The complete simulation chain with all steps described above, combined with the event reconstruction described in the next Chapter, are referred to as *full simulation*. As the full simulation requires several minutes of computation time for each event, it can be unfeasible in cases where a large number of simulated events is needed. Such cases include generation of numerous signal samples to scan a multi-dimensional parameter space of some theoretical model, as well as evaluation of modeling-related systematic uncertainties by generating several samples with different input parameters.

For these types of cases, an alternative approach known as *fast simulation* can be used. In fast simulation, a simplified detector geometry model is used, and instead of performing a detailed GEANT4 simulation in small steps, analytical models describing the particle interactions with the detector materials are applied. Their parameters are carefully tuned to match the full simulation. In the reconstruction step, the time-consuming iterative tracking (Section 8.1.1) is streamlined in fast simulation in two ways: the reconstructed hits are generated by simply smearing the simulated hits, and the generator-level information about the true particle trajectories is used to restrict the allowed hit combinations in the track finding step. As a result, the complete simulation and reconstruction chain is ≈ 20 times faster compared to the full simulation. Typically the results from fast and full simulations agree within ≈ 10 per cent. The fast simulation method is detailed in Ref. [154].

Chapter 8

Event reconstruction

The event reconstruction in CMS relies on a *particle-flow* (PF) algorithm [155]. Objects reconstructed with the PF are used in nearly all CMS analyses.

Traditionally, each type of particle has been reconstructed separately, combining different pieces of useful information from various subdetectors. The PF algorithm aims for a global event description, where a complete list of particles present in the event and their properties (charges, trajectories, and momenta) is produced by combining the full information from different subsystems in an optimized way. This PF algorithm concept is illustrated in Figure 8.1.

The PF approach was originally developed in the ALEPH experiment recording the e^+e^- collisions at LEP [157]. CMS was the first experiment to demonstrate that the PF approach can provide superior performance also in the reconstruction of the pp collisions, compared to the traditional methods. For the CMS, a key benefit of the PF method is that the energies (momenta) of charged hadrons are estimated more precisely by the highly segmented tracker, compared to less precise calorimeters. Also the ATLAS experiment has adopted the PF approach for some analyses [158].

The output of the PF algorithm is a set of *PF candidates*, classified into muons, electrons, photons, and charged and neutral hadrons. Ideally these correspond to the collection of stable particles present in the event. The PF candidates are then used to construct higher-level *physics objects*. Hadrons and photons are clustered into jets, and some jets can be further identified as originating from b quark hadronization or from hadronically decaying tau leptons. Finally, the presence of weakly interacting neutral particles can be inferred from the transverse momentum imbalance in the event by calculating the missing transverse momentum ($\vec{p}_T^{\text{miss}}$) from the reconstructed transverse momenta.
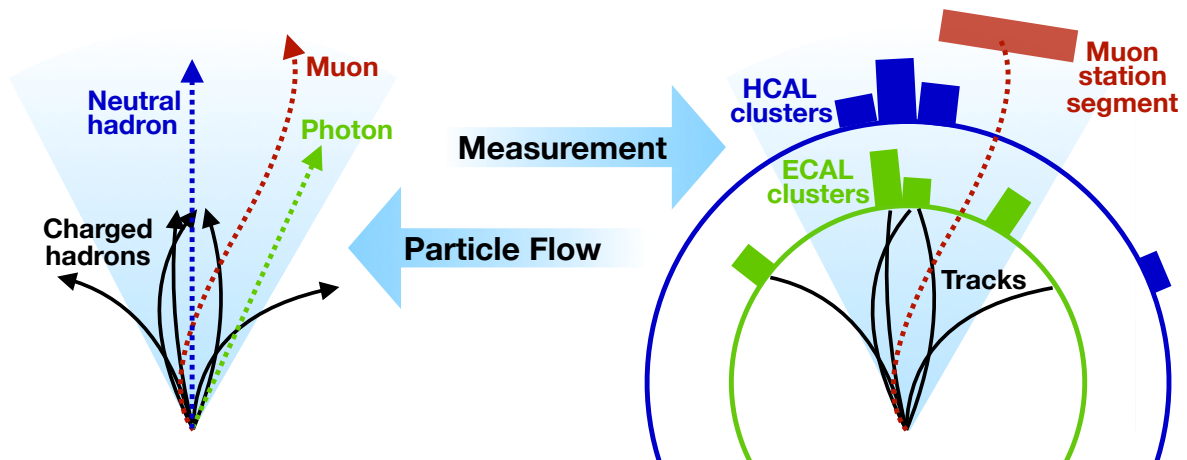
Figure 8.1: The Particle Flow algorithm concept. The PF method combines information from all subdetectors, to obtain a global view of the event, producing a list of measured particles that ideally correspond to the particles that were present in the real event. Illustration based on Ref. [156].

The reconstruction of each physics object relies on different parts of the detector: While jets are largely based on calorimeter deposits, the identification of τ_h candidates and b jets is significantly improved by the tracker information. Photons and electrons correspond to electromagnetic showers in the ECAL, and the muon chambers play a crucial role in muon identification. Fine granularity of subdetectors is needed to avoid merging different objects with close proximity and to reduce incorrect combinations of signals from different subdetectors.

The procedure to reconstruct the PF candidates is presented in Section 8.1. The physics objects constructed from the PF candidates and used in the analysis presented in this thesis, namely the isolated electrons and muons, jets, b jets, hadronically decaying tau leptons and missing transverse momentum, are described in Section 8.2.

8.1 Particle flow algorithm

Several steps are needed to reconstruct the PF candidates. First, the tracks and vertices are reconstructed from the hits in the tracker. Secondly, the calorimeter deposits are clustered and calibrated. Then tracks, calorimeter clusters and muon chamber signals are linked together and the resulting combinations are classified into different types of PF candidates.

8.1.1 Track reconstruction

The reconstruction of tracks from charged particles is performed iteratively with the combinatorial track finder (CTF) algorithm, based on the combinatorial Kalman filtering [159]. After identifying track seeds (a few hits compatible with a trajectory), pattern recognition is used to identify all hits along this trajectory. Then a fit to all hits is performed to determine the full track and its $\vec{p}_T$.

In the context of track reconstruction, the efficiency is defined as the fraction of simulated tracks successfully reconstructed such that at least half of the hits in the reconstructed tracks match with simulation. The fraction of simulated tracks failing this criterion defines the misreconstruction rate. The efficiency and misreconstruction rate for simulated QCD multijet events are shown in Figure 8.2. The black curve corresponds to a simple application of the CTF algorithm, reconstructing only high-quality tracks reconstructed from at least eight hits. We can see that the reconstruction efficiency reaches only 80% at its best. The main reason is that while the CMS multi-layer tracking detector gives high precision for the reconstructed tracks, it contains so much material that charged pions have 10–30% probability to undergo a nuclear interaction before reaching the first layers of the strip tracker, required for successful reconstruction [155].

The reconstruction efficiency can be improved without a large increase in misreconstruction rate by performing the track finding iteratively as described in Ref. [159]. The algorithm is applied in ten successive iterations using different seeding and quality criteria each time. The first iterations target the highest-quality tracks with clear signatures. After each iteration, the hits corresponding to reconstructed tracks are masked, so the combinatorial complexity is gradually reduced. Thus increasingly complex and computationally heavy criteria can be used in seeding, filtering and fitting as the iteration progresses. Some iterations are designed to tackle specific classes of tracks, such as tracks with two hits missing in the pixel detector, tracks from displaced vertices and tracks in dense high-p_T jets.

Figure 8.2 demonstrates the improvement in tracking performance from iterative tracking. The green curve corresponds to the subset of iterations with at least one seed hit in the pixel detector, the red curve corresponds to the full set of ten iterations, and the black curve shows the result from a single iteration. We can see that with ten iterations, a significant increase in efficiency is achieved with only a minor increase in the misreconstruction rate, as tracks affected by multiple scattering are partially

Figure 8.2: Efficiency (left) and misreconstruction rate (right) as a function of the track p_T for the CTF tracking algorithm after a single iteration (black squares), after multiple iterations with at least one seed hit in the pixel detector (green triangles) and after all iterations, including those with displaced track seeds (red circles). [155]

recovered. However, even with the gains from iterative tracking, a clear trend is visible: the efficiency decreases and the misreconstruction increases towards higher track p_T values. The reason is that most of the high-p_T tracks belong to dense collimated jets, containing multiple particles with trajectories close to each other, while the tracker resolution decreases with p_T. Therefore the resolution is not sufficient to disentangle them. This effect can be mitigated in the PF algorithm by linking the tracking information with the calorimeter information as discussed later in Section 8.1.4.

8.1.2 Vertex reconstruction

The collision vertices are reconstructed from particle tracks by extrapolating them from the tracker region towards the interaction point, and estimating the most probable vertex configuration using the deterministic annealing algorithm [160]. This technique results in approximately 10 μm resolution for the vertex locations in all three spatial dimensions.

In analyses targeting high-p_T objects, such as the one presented in this thesis, the primary pp interaction vertex is selected using so-called *track jets*. The track jets are formed by applying the anti-k_T jet clustering algorithm [161, 162] to the tracks associated with each vertex. After the track jets have been reconstructed, the associated

missing transverse momentum is determined as the negative vector sum of the $\vec{p}_T$ of the track jets. The p_T of each jet and the missing transverse momentum are summed quadratically. The vertex with the highest value of this p_T^2 sum is selected as the primary vertex.

8.1.3 Calorimeter deposit clustering

The electromagnetic showers in the ECAL and the hadron showers in the HCAL are wider than a single ECAL crystal or HCAL module. Thus clustering of the energy deposits in these crystals and modules, commonly referred to as *calorimeter cells*, is needed to determine the energies of the particles that initiate the showers.

Even though the clustering is performed separately in each calorimeter subsystem (ECAL and HCAL barrel and endcaps, and the two preshower layers in each endcap), the same approach is used everywhere: The clustering is seeded by cells where deposited energy exceeds a given energy threshold. The neighboring cells (8 closest in the ECAL and ES, 4 closest in the HCAL) are associated with the seeds. Then topological clustering is performed, adding nearby cells as long as the energy in additional cells is at least twice as large as the noise level. Finally, a maximum-likelihood fit based on a Gaussian-mixture model is used to reconstruct a set of Gaussian energy distributions inside each topological cluster, corresponding to electromagnetic or hadronic showers. Different seeding, clustering and fitting parameters are used in each subdetector, as listed in Ref. [155]. In HF, no clustering is used, and the energy deposits are interpreted directly as clusters.

For the purposes of electron and photon reconstruction, the ECAL energy from electromagnetic showers (initiated by an electron or a photon) is also summed into *superclusters* with a small width in η but a large coverage in ϕ to account for bremsstrahlung photons radiated by electrons, conversion of photons into electron–positron pairs, and the bending of the electron trajectories in the magnetic field. Superclusters are constructed by starting from a seed crystal, finding 5×1 (5×5) arrays around the seeds in the barrel (endcaps), and finally combining these into superclusters. Different minimum energy thresholds are applied for seeds, arrays and for the final supercluster. The supercluster positions are extrapolated to the ECAL preshower and the corresponding preshower energy deposits are added in the superclusters.

The response of the calorimeters is carefully calibrated. The initial calibration of ECAL was performed already before the start of LHC data taking using test beams,

radioactive sources and cosmic rays. During the operation of the detector, the changes in transparency of the ECAL crystals are monitored using a dedicated laser system, and the results are used to derive time-dependent corrections. These calibrations are complemented by a residual correction to the cluster energies, derived using GEANT4 simulation with the full CMS detector. The correction accounts for the effect of clustering thresholds (up to 20% corrections to cluster energies) and the shadowing of the ECAL endcap crystals by the preshower (corrections up to 40%). After these corrections, the calibrated energies and the simulated energies of photons typically agree at the percent level. The corrections are validated by studying data events with two energetic photons originating from π^0 decays, reconstructing the invariant mass of the diphoton system and comparing it to the nominal mass of π^0.

The ECAL supercluster energies are corrected and calibrated using an MVA regression algorithm trained with simulated events. The small residual differences between data and simulation are corrected by comparing the invariant mass peaks reconstructed in $Z \to e^+e^-$ events.

As the HCAL was initially calibrated with a test beam setup where the ECAL was not present, further calibration is needed to account for the fact that most hadrons deposit part of their energy already in the ECAL before reaching the HCAL. The calibrated calorimetric energy associated with a hadron is calculated from the calibrated ECAL and HCAL energies (E_{ECAL} and E_{HCAL}) as

$$E_{\text{calibrated}} = a + b(E)f(\eta)E_{\text{ECAL}} + c(E)g(\eta)E_{\text{HCAL}}, \quad (8.1)$$

where the constant coefficient a accounts for the energy missed due to the clustering thresholds, and the functions b/c (f/g) parametrize the ECAL/HCAL response as a function of hadrons simulated energy E (pseudorapidity η). Events with neutral hadrons processed through a GEANT4 detector simulation are used to derive the calibration coefficients. They are derived separately for different calorimeter regions and separately for showers initiated in the ECAL and those contained only in the HCAL. The calibration is found to yield essentially linear energy response and improved energy resolution, and it has been validated using real collision events with isolated charged hadrons. The details of the calibration and its validation are given in Ref. [155].

8.1.4 Link algorithm

Once the tracks have been reconstructed with iterative tracking, a *link algorithm* is used to match them with the calibrated calorimeter clusters and signals in muon chambers.

The links between the different *PF elements* (tracks, clusters, superclusters and muon chamber signals) are established by extrapolating the trajectories reconstructed in the tracker, taking into account the effect of the magnetic field. The elements are matched geometrically in the (η, ϕ) plane. For example, if the extrapolated track passes through a cell which is part of a HCAL cluster, the cluster is linked to the track. If multiple combinations of tracks and clusters are possible, the one with the minimal ΔR distance is chosen.

As bremsstrahlung can cause sudden changes in the curvature of electron trajectories, the standard CTF method can lead to lost hits and poor fit quality. Thus the electron track candidates (tracks linked with ECAL superclusters) are refitted with a dedicated Gaussian-sum filter (GSF) algorithm that takes into account energy losses due to bremsstrahlung in the tracker. Photons from electron bremsstrahlung are linked by extrapolating the tangents of GSF tracks to ECAL. A dedicated algorithm is used to find tracks from e^+e^- pairs compatible with originating from a photon conversion. Such tracks are linked with the original track of the electron radiating the photon.

The ECAL, HCAL and ES clusters as well as ECAL superclusters are also linked with each other by geometrical matching. Tracks originating from the same secondary interaction vertex are linked together. Finally, tracks found in the tracker and signals in the muon detectors are linked together for muon identification.

The output of the link algorithm is a set of *PF blocks*, each containing a collection of elements linked to each other, typically originating from one or only a few particles.

8.1.5 Reconstruction of the particle-flow candidates

The identification and reconstruction of particle-flow candidates is carried out in each PF block separately. The muons are reconstructed first (tracker and muon chamber tracks linked), followed by electrons (tracks and ECAL clusters linked) and isolated photons (isolated ECAL clusters). After reconstruction of each type of PF candidate, the corresponding tracks and clusters are removed from the PF block. The hadrons and non-isolated photons, arising from fragmentation and hadronization in jets, are reconstructed last. The remaining ECAL+HCAL clusters with (without) a linked track

are identified as charged (neutral) hadrons, and the ECAL clusters without a linked track as nonisolated photons. Outside the tracker coverage, the ECAL+HCAL deposits are classified as "hadronic energy" and the isolated ECAL deposits as "electromagnetic energy", without associating them to any particles. Details on the reconstruction of each type of PF candidate are given in the following.

Muons:

As calorimeters absorb nearly all particles except muons and neutrinos, muons can be reconstructed with high purity based on signals in muon chambers. As discussed in Section 5.5, signals from the muon chambers can be reconstructed as *hits* with well-defined positions. While the RPC chambers contain only a single layer, the CSC and DT chambers have multiple layers, so hits in their subsequent layers are connected by fitting a straight line through them to form *segments*. Three types of muons can be reconstructed: standalone muons, global muons and track muons [163]. The joint collection of these different types of muons is referred to as *PF muons*. According to simulation studies, the muon reconstruction efficiency for PF muons is > 99%.

The *standalone muons* rely only on the muon chamber signals. They are built using the Kalman filtering method [164], by choosing the DT and CSC segments as seeds, and using pattern recognition to find associated hits in all muon detectors (CSC, DT and RPC). Then a track fit is performed using all these hits.

The *global muons* contain a standalone muon track matched to another track, reconstructed in the tracker system as described in Section 8.1.1. The matching is performed by extrapolating both tracks to a common surface. After matching, a global track fit is performed with the Kalman filter method.

The *tracker muons* are built by extrapolating reconstructed tracks from tracker to muon systems in the transverse plane, and requiring that at least one DT or CSC segment matches to the extrapolated track. The tracker muons complement the other approaches especially in the $p_T < 10\,\text{GeV}$ range, where it is common that a proper track cannot be reconstructed in muon chambers due to multiple scattering in the magnetic return yoke, and typically only one matching segment in the innermost muon stations is found. The muons need to have a minimum three-momentum of a few GeV to reach even the innermost muon stations. For muon with (almost) transverse trajectories, this corresponds to a few GeV, while for muons with a large longitudinal momentum component the p_T threshold is even lower.

In the PF algorithm, first the isolated global muons are identified in the given PF block. The isolation is defined by requiring the sum of the p_T of the tracks and the E_T of the calorimeter energy deposits to be less than 10% of the muon p_T. For nonisolated muons (typically found inside jets), tight selection criteria as described in Section 8.2.1 are required. Additionally, three track segments in muon chambers or lack of significant calorimeter deposits is required to suppress punch-through hadrons misidentified as muons. As the tight selection criteria require successful muon track reconstruction in both the muon chambers and the tracker, also muon candidates with a large number of hits in the tracker alone (tracker muons) or muon chambers alone (standalone muons) are accepted. The details are given in Ref. [155].

The momenta of muons are obtained from the curvatures of the corresponding tracks. The tracker system provides the best momentum resolution up to $p_T \approx 200\,\text{GeV}$, whereas for higher p_T values the inclusion of muon chamber track curvature improves the estimate. To optimize the momentum resolution across the momentum range, for each muon candidate above $200\,\text{GeV}$ the p_T is calculated using several combinations of tracker and muon information, and the most precise result is chosen.

Electrons and isolated photons:

As electrons pass through the tracker, they are likely to radiate bremsstrahlung photons. According to simulation studies, on average 33% (86%) of the electron energy is radiated for $|\eta| \approx 0$ ($|\eta| \approx 1.4$), corresponding to the smallest (largest) amount of tracker material to traverse before reaching ECAL. Both prompt and bremsstrahlung photons, on the other hand, are likely to convert to e^+e^- pairs, which then radiate photons. Thus it is convenient to reconstruct electrons and isolated photons using a common approach.

While isolated electrons with large enough momenta are typically associated with ECAL superclusters, this is not always the case for non-isolated electrons inside jets or small-p_T electrons with large track bending. Thus the PF electron candidates are seeded either by an energy deposit in the ECAL (supercluster $E_T > 4\,\text{GeV}$) or by reconstructed tracks with $p_T > 2\,\text{GeV}$, linked with an ECAL supercluster. A BDT that combines several variables describing the properties of calorimeter clusters, as well as the direction and quality of the associated tracks, is used to identify the electrons as described in Section 8.2.2.

Isolated photons are seeded by ECAL superclusters with $E_T > 10\,\text{GeV}$ not matched to a GSF track. The candidates are required to be isolated both based on reconstructed

tracks and calorimeter clusters. The ratio of the ECAL and HCAL energy deposits is required to be compatible with an electromagnetic shower initiated by a photon. For the candidates passing these requirements, the direction and the corrected energy of the supercluster are taken as the PF photon energy and direction.

Hadrons and non-isolated photons:

Once muons, electrons and isolated photons have been reconstructed and the PF elements associated to them removed from the PF blocks, hadrons produced in fragmentation and hadronization of quarks and gluons can be reconstructed. As e.g. neutral pions decay to photons, also the nonisolated photons are identified at this point.

All remaining ECAL clusters (with or without a linked HCAL cluster) within the tracker acceptance ($|\eta| < 2.5$) but without associated tracks are interpreted as nonisolated photons. Similarly, all HCAL clusters without associated tracks are interpreted as neutral hadrons.

Outside the tracker acceptance the charged and neutral hadrons cannot be separated, so the classification is based on calorimeter information alone: the ECAL clusters with no linked HCAL clusters are interpreted as photons, while all HCAL clusters (with or without a linked ECAL cluster) are classified as hadrons.

For the calorimeter clusters associated with tracks, the sum of track momenta is compared to the calibrated cluster energy. Each track is interpreted as a charged pion, the momentum of which is recalculated with a combined fit in the associated track and calorimeter deposits. If the clustered calorimetric energy is larger than the track momenta, the additional ECAL energy is associated with photons and HCAL energy with neutral hadrons.

Finally, in cases where calorimetric energy is smaller than the sum of track momenta, additional global muons earlier obscured by the presence of other particles are identified and added to the list of PF muons. Anomalous cases from misreconstructed tracks are solved by masking tracks with large p_T uncertainty. Charged-particle tracks associated with secondary vertices are used to reconstruct the original charged hadrons that were present before nuclear interactions took place. The reconstructed hadrons are associated with the mass of a charged pion and used to replace their decay products in the list of final PF objects.

Missing transverse energy:

The missing transverse momentum $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ is defined as the negative vector sum of the $\vec{p}_{\mathrm{T}}$ of all reconstructed PF candidates [165]. The reconstructed events can contain genuine $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ from neutrinos (or from other weakly interacting particles not yet known to us) that escape the detector undetected, but also *anomalous* $p_{\mathrm{T}}^{\mathrm{miss}}$ originating from misidentification of particles and mismeasurement of particle momenta.

As large $p_{\mathrm{T}}^{\mathrm{miss}}$ is an important signature for specific SM processes as well as many BSM signals, care must be taken to suppress any anomalous $p_{\mathrm{T}}^{\mathrm{miss}}$ as well as possible. Thus different types of high-p_{T} particles that might produce artificial $p_{\mathrm{T}}^{\mathrm{miss}}$ are identified and the correlation between their p_{T} and the $p_{\mathrm{T}}^{\mathrm{miss}}$ is studied. If a large correlation is found, the choices made during the PF reconstruction are re-evaluated and an alternative interpretation of the event that reduces the $p_{\mathrm{T}}^{\mathrm{miss}}$ is sought for.

Firstly, muons from cosmic rays can produce anomalous $p_{\mathrm{T}}^{\mathrm{miss}}$. If muons are found with trajectories more than 1 cm away from the beam axis and their presence affects the $p_{\mathrm{T}}^{\mathrm{miss}}$ significantly, they are removed (the latter condition is required in order to protect genuine displaced muons e.g. from b hadron decays).

A second source of anomalous $p_{\mathrm{T}}^{\mathrm{miss}}$ is the misreconstruction of muon momenta due to e.g. wrong track association, decay or interaction in the return yoke. Thus for all muons with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$, different momentum reconstruction methods discussed in Section 8.1.5 are studied and the one yielding the lowest $p_{\mathrm{T}}^{\mathrm{miss}}$ is chosen.

Thirdly, for muons with $p_{\mathrm{T}} > 100\,\mathrm{GeV}$ associated with neutral hadron energy above $100\,\mathrm{GeV}$, the punch-through hadron hypothesis is studied. Similarly, charged hadrons almost collinear with $p_{\mathrm{T}}^{\mathrm{miss}}$ are tested for a hypothesis that they are actually misidentified muons and $p_{\mathrm{T}}^{\mathrm{miss}}$ originates from a nonreconstructed neutral hadron (almost collinear with the charged hadron). In both cases, the alternative hypothesis is accepted and the list of PF candidates modified accordingly if the new hypothesis reduces $p_{\mathrm{T}}^{\mathrm{miss}}$ by at least 50%.

As the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ can be calculated only after all PF candidates are reconstructed, these corrections are performed at the end of the PF candidate reconstruction sequence. Simulation studies confirm that events with genuine $p_{\mathrm{T}}^{\mathrm{miss}}$ are not significantly modified by the corrections [165].

8.2 Reconstruction of physics objects

8.2.1 Isolated muons

As discussed in Section 8.1.5, the PF muons are reconstructed and their momentum is precisely determined by combining information from the tracker and muon chambers in an optimized way.

Based on the properties of the PF muons, different identification criteria are defined. Any PF candidate classified as a muon and passing the selection criteria for tracker muons or global muons is defined to pass the *loose muon identification*. If a loose muon candidate is reconstructed both as a tracker muon and a global muon, has a tracker track with associated hits in at least six layers of the tracker and with at least one of them in the pixel detector, has a segment traversing at least two muon stations, and has a high-quality global track with $\chi^2/N_{\mathrm{dof}} < 10$ compatible with the primary vertex and with transverse (longitudinal) impact parameter smaller than 0.2 (0.5) cm, it is defined to pass the *tight muon identification*. The aim of these stringent criteria for tight muons is to select a pure sample of prompt muons, efficiently suppressing secondary muons from decays in flight, misidentified muons originating from punch-through and cosmic muons. Other identification criteria have been developed to improve the selection efficiency specifically for e.g. low-p_T or high-p_T muons and they are described in Ref. [163].

To discriminate prompt, isolated muons from secondary muons present in jets (originating e.g. from b hadron decays), the isolation criteria are often defined based on reconstructed tracks or PF candidates present in the proximity of the muon candidate. One commonly used isolation discriminant is the so-called *mini-isolation*, defined as the p_T sum of the PF candidates inside a cone around the muon. For optimal performance across the muon p_T range, the cone size is varied as a function of p_T. For $p_\mathrm{T} < 50\,\mathrm{GeV}$, the cone radius is 0.2. For $50 < p_\mathrm{T} < 200\,\mathrm{GeV}$, the cone radius is $\Delta R = 10.0\,\mathrm{GeV}/p_\mathrm{T}$, and for $p_\mathrm{T} > 200\,\mathrm{GeV}$ it is 0.05.

8.2.2 Isolated electrons

As mentioned in Section 8.1.5, the electron reconstruction can be seeded in two complementary ways. Most electrons can be reconstructed from seeds based on ECAL superclusters: the supercluster position is calculated as the energy-weighted average and back-propagated to the tracker, assuming no radiation losses and considering both

negative and positive charge hypotheses. The electron candidate is built by associating the matching CTF and GSF tracks with the supercluster. Only superclusters with $E_{\mathrm{T}} > 4\,\mathrm{GeV}$ and no significant matching energy deposits in HCAL are considered. Alternatively, electron reconstruction can be seeded by a reconstructed CTF track with $p_{\mathrm{T}} > 2\,\mathrm{GeV}$, and linked with an ECAL cluster. In this case, the additional ECAL deposits from bremsstrahlung are included by extrapolating tangents of the CTF track to ECAL and associating the found ECAL clusters with the candidate. The track seeds improve reconstruction efficiency for low-p_{T} electrons as well as electrons falling in the ECAL barrel-endcap transition region.

The energies of the ECAL clusters are corrected as a function of the cluster energy and η to compensate for the energy lost in the tracker. The electron candidates are required to have their cluster energy and GSF track momentum of similar magnitude.

Finally, a BDT combining information from the tracks and the associated calorimeter clusters is used to identify the electrons. The BDT considers several properties of the candidate such as the CTF and GSF track fits and their difference, the compatibility of the candidate direction and momentum as estimated from tracks and clusters, supercluster substructure, and the fraction of energy lost in bremsstrahlung (estimated by comparing the reconstructed momentum in the center of the tracker and in the ECAL surface). The BDT is trained in bins of electron p_{T} and η, separately for barrel and endcap regions and for isolated and nonisolated electrons.

For candidates passing the BDT selection a second, more detailed GSF track fit is performed to determine the $\vec{p}_{\mathrm{T}}$ of the electron at high precision. The final momentum estimate is then obtained by combining the track momentum and the associated supercluster energy as a weighted average. The momentum resolution for electrons with $p_{\mathrm{T}} \approx 45\,\mathrm{GeV}$ from $Z \to \mathrm{ee}$ decays ranges from 1.7% for nonshowering electrons in the barrel region to 4.5% for showering electrons in the endcaps [166].

To separate prompt electrons from jets misidentified as electrons and from nonprompt electrons produced within a jet (e.g. in b hadron decays), electrons are typically required to be isolated from hadronic activity in the event. The isolation variable can be based on the calorimeter deposits, tracks or PF candidates, with their energies summed inside an isolation cone that is defined around the electron candidate. Similarly to muons, a *mini-isolation* discriminator is defined for electrons as the p_{T} sum of PF candidates inside a p_{T}-dependent cone around the electron, with cone size $\Delta R = 10\,\mathrm{GeV}/\min(\max(p_{\mathrm{T}}, 50\,\mathrm{GeV}), 200\,\mathrm{GeV})$.

8.2.3 Jets

For each event, hadronic jets are clustered from the reconstructed PF candidates using the anti-k_T algorithm [161, 162] with a distance parameter $R = 0.4$.

The anti-k_T algorithm starts from a seed particle (PF candidate) i. For each surrounding particle j, a distance measured in terms of *rapidity y* and angle ϕ is calculated as $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$. Then, a variable d_{ij} that reflects both the distance and the transverse momenta of the particles is determined as $d_{ij} = \min(p_{T_i}^{-2}, p_{T_j}^{-2})\Delta_{ij}^2/R^2$. First the particle j that minimizes d_{ij} is chosen, and if $d_{ij} < p_{T_i}^{-2}$, the particle j is clustered into the same jet with the particle i. The clustering continues by considering another particle j (not yet clustered into the jet) that minimizes d_{ij}, and continues until all particles with $d_{ij} < p_{T_i}^{-2}$ have been clustered to the jet containing the particle i. Then the algorithm proceeds to build the next jet around a different seed particle i.

The chosen distance measure d_{ij} ensures that soft (i.e. low-momentum) particles tend to cluster into jets with hard (i.e. high-momentum) particles rather than among themselves, so the anti-k_T algorithm tends to produce a collection of hard jets [161]. The anti-k_T algorithm is *infrared safe*, so that addition of very soft ("infrared") particles does not change the outcome of the clustering. It is also *collinear safe*, meaning that a set of collinear particles with some total momentum are clustered in the same way as a single particle with this momentum.

After clustering, the jet momentum is determined as the vectorial sum of all particle momenta in the jet, and is found from simulation to be within 5 to 10% of the true momentum over the whole p_T spectrum and detector acceptance.

Pileup can contribute additional tracks and calorimetric energy deposits to the jet momentum. To mitigate this effect, the *charged hadron subtraction* algorithm [167] is used to identify and remove the charged hadron tracks associated to pileup vertices. These charged hadrons correspond to approximately two thirds of the energy deposited due to pileup. The remaining neutral hadron contribution is corrected as part of the following step, the jet calibration.

In this thesis, the charged hadron subtraction technique is used. An alternative technique for pileup mitigation, known a *pileup per particle identification* [168], is also used in some CMS analyses. In this approach, each reconstructed particle is assigned a weight that reflects the probability that the particle originates from a pileup interaction, and the weights are used to scale the particle four-momenta accordingly.

Figure 8.3: Workflow of jet energy calibration. Pileup offset corrections and simulated response corrections are applied to both data and simulated events, whereas the last steps of residual corrections (based on γ +jet, Z +jet and dijet events) are applied only to data. [169]

Jet calibration:

Despite careful calibration of calorimeter clusters and precise charged-particle momenta provided by the tracker, further calibration is needed to make the momenta of the clustered PF jets match with the momenta of corresponding particle-level jets. Particle-level jets are defined based on generator-level information in simulated events, clustering all stable particles (decay length $c\tau > 1$ cm) except neutrinos into jets with the anti-k_T algorithm. Differences between reconstructed and particle-level jets arise from several sources: the energy of the reconstructed jets is affected by energy deposits from pileup interactions (the *pileup offset*), initial and final state radiation, electronics noise, and other detector effects.

The jet energy calibration is factorized into several steps (levels): pileup offset corrections, simulated response corrections, and residual corrections. The first steps are applied to both data and simulated events, while the residual corrections are only applied to data to bring them closer to the simulation. Each correction is applied as a multiplicative correction factor on the jet four-momentum. These steps are illustrated in Figure 8.3. A detailed description of each correction step is given in Ref. [167].

The pileup offset corrections are designed to correct the jet energy by estimating and subtracting the energy corresponding to pileup inside a jet. For this purpose, correction factors are first derived from simulation by reconstructing the same simulated QCD multijet events with and without pileup simulation included, and comparing the resulting reconstructed jets. The corrections are applied as a function of the (uncorrected) jet p_T, η, and the *jet area*, defined as the (η, ϕ) region where the particles get clustered into the jet. The correction also depends on the *diffuse offset energy density ρ* of the event, used to parametrize the pileup dependence. It is estimated as a median of energies deposited in η/ϕ bins covering the full detector [167].

To cover for difference in pileup offset between data and simulation, the *random-cone method* is used to derive additional corrections. This method is used with so-called zero-bias events, which are triggered completely randomly (only requiring that a bunch crossing between two filled bunches take place) so that most of them contain only pileup vertices and no hard scatterings. The energy deposition distributions between data and simulation are compared by summing the PF candidate energies inside cones placed at random η and ϕ coordinates in both simulated and reconstructed zero-bias events. The resulting correction is applied to simulated events as a function of η and ρ. The correction factor is typically of the order of 10%, but can be significantly larger or smaller in specific bins [169].

After the pileup offset correction, the jet simulated response corrections are applied. The momentum response is defined as $\langle p_{\mathrm{T}} \rangle / \langle p_{\mathrm{T}}^{\mathrm{gen}} \rangle$ where $\langle p_{\mathrm{T}} \rangle$ is the mean p_{T} of the reconstructed jets in a given $p_{\mathrm{T}}^{\mathrm{gen}}$ bin, and $\langle p_{\mathrm{T}}^{\mathrm{gen}} \rangle$ is the mean of the generator-level jet transverse momenta in the same bin. The reconstructed jets are matched to the generator-level jets by requiring the distance between two jets to be $\Delta R < 0.2$. The corrections vary considerably between different regions of the detector, they are applied in bins of η (as a function of the reconstructed jet p_{T}). These corrections are typically at the percent level, although they can be larger in specific p_{T}/η regions [169].

After pileup offset and simulated response corrections, the remaining differences between simulation and data are mitigated with a series of residual corrections. The idea is to take advantage of momentum conservation in the transverse plane, and use an object with a precisely measured energy scale (a Z boson, a photon or another jet) to estimate the energy scale of a recoiling jet. Firstly, information from the precise energy measurements in the barrel region is extrapolated to endcap regions by exploiting dijet events that contain two jets with very similar p_{T}, one in the barrel ($|\eta| < 1.3$) and the other outside it. Secondly, $Z \to e^{+}e^{-}$+jet, $Z \to \mu^{+}\mu^{-}$+jet, and γ+jet events are used to correct the jet momenta in the p_{T} range of 30–700 GeV. The p_{T} of the Z boson or photon is measured at high precision and then used to estimate the miscalibration and derive a correction for the recoiling jet. Thirdly, the calibrations can be extrapolated to the TeV range by using QCD multijet events with one high-p_{T} jet and ≥ 2 recoiling softer jets. A simultaneous fit combining these different types of events is performed, determining the absolute jet p_{T} scale via chi-squared minimization. The energy scales of other physics objects are allowed to vary within their uncertainties.

Figure 8.4: Illustration of the flight distance defined as the distance between the primary vertex (PV) and the secondary vertex (SV), and the impact parameter (IP) defined as the distance from the PV to the track at its closest point of approach. [171]

After all corrections, the jet energy resolution amounts typically to 15% at 10 GeV, 8% at 100 GeV, and 4% at 1 TeV [170]. Each level of corrections is assigned systematic uncertainties as described in Ref. [167].

8.2.4 Heavy-flavor jets

Identification of the jets originating from the hadronization of bottom quarks (b jets) or charm quarks (c jets) is necessary to efficiently select events with these quarks in the final state. Several algorithms are used in CMS to perform this *b-tagging* (or *c-tagging*).

The hadronization of b and c quarks leads to the formation of short-lived heavy-flavor hadrons, with lifetimes of the order of 1 ps or less. Depending on the hadron p_T, this corresponds to flight distances from a few mm up to 1 cm. The decay of these hadrons is observed as secondary vertices that can be reconstructed from the tracks of the decay products. The fact that the b and c quarks are heavy compared to the light (u,d,s) quarks and massless gluons affects the kinematic properties of the jet constituents, such as the distribution of the p_T relative to the jet axis. Also the presence of *soft leptons*, i.e. electrons and muons from heavy flavor hadron decays can be used for identification. While they are present only in 20% (10%) of the b (c) jets, they allow a selection of a pure sample of heavy flavor jets.

In the simulated events, the b and c jets are defined as follows. The jets containing at least one b hadron are classified as b jets, while jets with at least one c hadron
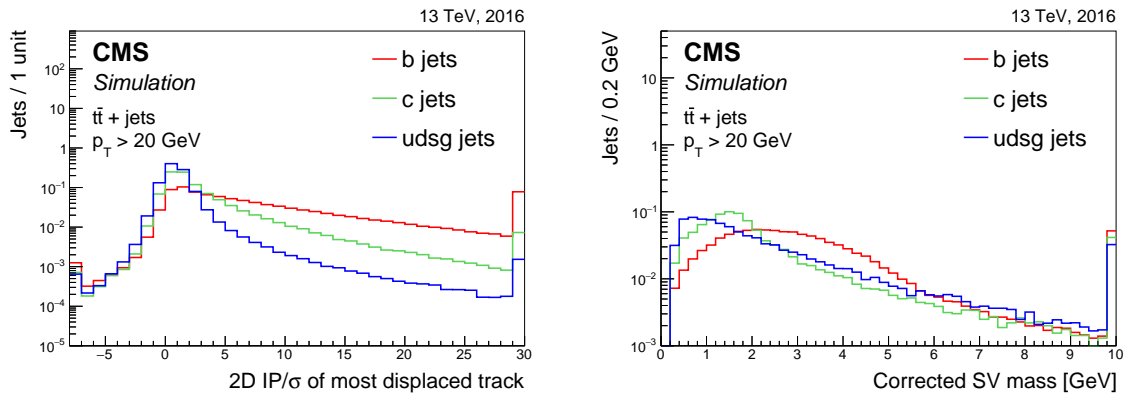
Figure 8.5: The distributions of the 2D impact parameter significance (left) and the reconstructed secondary-vertex mass (right) for different types of jets with $p_T > 20\,\text{GeV}$ in simulated $t\bar{t}$ events. [171]

and no b hadrons are labeled c jets. The generator-level hadrons are assigned to the reconstructed jets by *ghost association* [172]. In this technique, the generator-level b and c hadrons are associated with jets by including them in the anti-k_T clustering of the PF candidates, but their masses are set to zero so that the jet momentum is not affected.

Several algorithms are used in CMS to perform the tagging of heavy flavor jets. The algorithms vary in complexity and use as input the properties of tracks, secondary vertices, soft leptons or a combination of these. A collection of high-quality tracks with $p_T > 1\,\text{GeV}$ is used as input, with additional selection criteria designed to suppress tracks from pileup, from long-lived K_S^0 and Λ hadrons and from misreconstruction. The secondary vertices (SV) are reconstructed using the inclusive vertex finding algorithm, providing reconstruction efficiencies of 75%, 37% and 12% for b, c, and light (u, d, s, g) jets, respectively, with jet $p_T > 20\,\text{GeV}$ in simulated $t\bar{t}$ events. The SV finding allows the reconstruction of the invariant mass of the hadron from the tracks associated to the SV, and the determination of the flight distance, defined as the distance to the primary vertex.

The impact parameter (IP) of a track is defined as the distance from the primary vertex to the track at its closest point of approach. It can be defined in three dimensions, in the transverse plane ("2D IP"), or along the z axis (longitudinal IP). The IP significance is defined as the ratio of the IP and its uncertainty. The definitions of the flight distance and of the impact parameter are illustrated in Figure 8.4. The distributions of the 2D IP significance of the most displaced track and the reconstructed SV mass are shown in Figure 8.5. Some algorithms also use as input the properties of soft leptons, i.e.

nonisolated electrons or muons clustered into jets and passing loose identification criteria, with $p_T > 2\,\text{GeV}$.

Jet probability tagger algorithms rely only on the IP significances of the tracks to determine the likelihood for a jet to originate from the PV. The *combined secondary vertex* (CSV) algorithm [173] makes use of track information, such as the IP significances and the track p_T and η relative to the jet axis, as well as the SV information (such as reconstructed masses of the SVs and the flight distance significances). The CSV algorithm updated and optimized for the LHC Run 2 is referred to as CSVv2, and uses 19 input variables, combining them into a multivariate discriminant with a shallow neural network. Another version of the CSV algorithm uses a deep neural network with more hidden layers and is labeled DeepCSV. Finally, a BDT-based combined tagger algorithm called cMVAv2 combines the outputs of several variants of jet probability tagger and CSV algorithms. The detailed descriptions of the different algorithms and studies on their performance can be found in Ref. [171].

8.2.5 Hadronic tau leptons

The τ_h are reconstructed with the hadron-plus-strips algorithm [174, 175], The algorithm is seeded by PF jets with $p_T > 14\,\text{GeV}$ (prior to the jet energy corrections), within the tracker acceptance ($|\eta| < 2.5$). As discussed in Section 2.5.1, most tau decays contain neutral pions that mostly decay into two photons. The photons can convert into e^+e^- pairs. As electrons are bent in the magnetic field, they radiate bremsstrahlung photons. To identify the jets corresponding to a τ_h, the HPS algorithm attempts to find the charged pions present in the jet and to combine them with *strips* of e/γ candidates corresponding to the photons from the neutral pion decay(s). The jets where a combination of hadrons and strips compatible with a hadronic tau decay is found are accepted as τ_h candidates.

The electron and photon candidates with $p_T > 0.5\,\text{GeV}$ are clustered into strips via an iterative process: The highest-p_T e/γ candidate in the jet is used to seed the clustering. A window rectangular in η and ϕ is defined around the candidate. Then the e/γ candidate with second-highest p_T is added into the strip, and the window is moved by calculating the p_T-weighted average of the η and ϕ coordinates of the e/γ candidates and placing the window around the averaged coordinates. Again, the highest-p_T candidate found inside the window and not yet clustered is added into the strip, and the process is continues until no e/γ candidates are left to add within the window.

Figure 8.6: Illustration of the HPS algorithm, using a $\tau^- \rightarrow \pi^- \pi^0 \nu_\tau$ decay as an example. On the left, the decay chain is shown schematically (the neutrino is not drawn). In the tracker, a photon is converted into opposite-signed electrons, bent in opposite directions by the magnetic field. The dotted lines represent the smaller strip clustering windows used during Run 1, and a larger window allowed by the Run-2 version of the HPS algorithm.

After one strip is finished, the algorithm proceeds to pick the highest-p_T e/γ candidate in the jet, and not yet included in any strip, as the new seed.

In the original version of the HPS algorithm, used during Run 1 [174], a constant strip size of $\Delta\eta \times \Delta\phi = 0.05 \times 0.20$ was used. In the updated Run 2 version of the algorithm [175], the window size is allowed to vary from 0.05×0.05 to 0.15×0.30 as

$$
\begin{aligned}
\Delta\eta &= f(p_T^{e/\gamma}) + f(p_T^{strip}), \\
\Delta\phi &= g(p_T^{e/\gamma}) + g(p_T^{strip}),
\end{aligned}
\tag{8.2}
$$

where $f(p_T) = 0.20 \cdot p_T^{-0.66}$ and $g(p_T) = 0.35 \cdot p_T^{-0.71}$. The exponentially decreasing functional form is motivated by simulation studies indicating that the e/γ showers from high-p_T tau lepton decays are more collimated in η and ϕ, while the decay products of lower-p_T tau leptons are more likely to undergo interactions with the tracker material, producing e/γ particles with larger separation from the original τ. The optimal numerical coefficients in the functions f and g are obtained from simulation [175].

One example of the τ_h decay signatures targeted by the HPS algorithm is shown in Figure 8.6. A tau lepton decays via $\tau^- \rightarrow \pi^- \pi^0 \nu_\tau$, the π^0 decays promptly to photons,

and one of the photons is converted into electrons in the tracker, as illustrated on the left (the neutrino is not drawn). On the right, these particles are shown as part of the jet (represented by a gray circle) in the (η, ϕ) plane. The dotted lines represent the smaller strip clustering windows used during Run 1, and a larger window allowed by the Run-2 version of the HPS algorithm.

To combine the strips with the charged hadrons present in the jet, the vertex closest to the highest-p_T charged particle in the jet is selected and considered as a possible τ decay vertex. All charged PF candidates with $p_T > 0.5\,\text{GeV}$, $\Delta z < 0.4\,\text{cm}$ and $\Delta r < 0.03\,\text{cm}$ are considered, where Δz (Δr) is the distance of closest approach between the vertex and the track along the beam axis (in the transverse plane).

After the strips and charged particles are found, different hypotheses corresponding to different tau decay modes are tested, considering the highest-p_T charged particles and strips, at most six of each. The invariant mass resulting from the combination is calculated by summing the four-momenta of the charged particles and of the e/γ candidates in the strips. For the mass calculation, the electrons and muon candidates are set to nominal e/μ masses, except for the (secondary) electrons associated with strips which are treated as massless, while all other charged particles are assigned the $\pi^\pm$ mass. The different combinations that are tested, the corresponding τ decay modes targeted, and the allowed range for the invariant mass of the combination are listed in Table 8.1. The decay mode $\tau^- \to h^- h^+ h^- \pi^0 \nu_\tau$ is not considered, because it is hard to distinguish from quark and gluon jets and has a small branching fraction of 4.8%.

The combinations that satisfy one of the decay mode hypotheses, pass the corresponding mass window selection, contain charges summing up to a ± 1 unit charge, and have no extra tracks or strips outside the ΔR cone of $3.0/p_T[\text{GeV}]$, varying from 0.05 to 0.1, are accepted as τ_h candidates. If several hypotheses using same particles or strips pass the selection, the candidate with the highest p_T is chosen.

When the HPS algorithm is applied in data, the resulting collection of τ_h candidates is contaminated by quark and gluon jets. The jet $\to \tau_h$ misidentification rate can be suppressed by exploiting the fact that the decay products of a hadronic tau decay are relatively isolated compared to quark and gluon jets. The τ_h isolation is calculated by summing the transverse momenta of all charged particles and photons not part of the τ_h candidate and with $p_T > 0.5\,\text{GeV}$ within a ΔR cone of (typically) 0.5.

Table 8.1: Different hadronic decay modes of tau leptons considered by the HPS algorithm and the corresponding selection criteria. Charge conjugation is implied.

Decay mode(s)	Combination	Min. m_{τ_h}	Max. m_{τ_h}	Comment
$\tau^- \to h^- \nu_\tau$	h^-	—	—	No associated strips
$\tau^- \to h^- \pi^0 \nu_\tau$	$h^- +$ strip	0.3	$1.3\sqrt{\frac{p_T[\text{GeV}]}{100\,\text{GeV}}}$	Max. m_{τ_h} within 1.3–4.2 GeV
$\tau^- \to h^- \pi^0 \pi^0 \nu_\tau$	$h^- +$ 2 strips	0.4	$1.2\sqrt{\frac{p_T[\text{GeV}]}{100\,\text{GeV}}}$	Max. m_{τ_h} within 1.2–4.0 GeV
$\tau^- \to h^- h^+ h^- \pi^0 \nu_\tau$	$h^- h^+ h^-$	0.8	1.5	All $h^\pm$ associated to same vertex

To discard charged hadrons from pileup, tracks of the charged particles are required to be compatible with the τ_h vertex ($\Delta z < 0.2$ cm, $\Delta r < 0.03$ cm). As photons do not leave tracks, the pileup photons cannot be avoided in a similar way. Instead, the contribution from pileup photons is estimated indirectly and subtracted from the photon p_T sum. First the pileup contribution from charged hadrons is calculated by summing the p_T of the charged particles with tracks failing the Δz criterion within a larger ΔR cone of radius 0.8. This sum ($\sum_{\text{charged,PU}}$) is then multiplied with so-called $\Delta\beta$ factor defined as the ratio of the neutral and charged components of the pileup. Thus the $\Delta\beta$ corrected isolation I is defined as

$$I = \sum_{\text{charged}} p_T + \max\left(0, \sum_\gamma p_T - \Delta\beta \sum_{\text{charged,PU}} p_T\right). \tag{8.3}$$

The value of $\Delta\beta = 0.2$, obtained from simulation studies, is used in Run 2 [175].

In order to discriminate between τ_h candidates from genuine tau leptons and jets misidentified as τ_h, the isolation information can be combined with information about the associated tracks via a dedicated BDT. The input parameters for the discriminant are: both terms of equation 8.3 (separately), the reconstructed τ_h decay mode, the transverse impact parameter of the leading track (defined as the distance of closest approach between the track and the associated vertex), the significance of the impact parameter (defined as the value divided by its uncertainty), the distance between the τ production and decay vertices as reconstructed with the adaptive vertex fitter algorithm [176] (if the reconstruction of the decay vertex was successful), and the significance of the vertex distance. Finally, the τ_h candidate p_T and $|\eta|$ are included in the inputs to keep track of how the other input variables depend on these quantities.

Figure 8.7: Performance curve of the BDT used for τ_h identification, trained separately for 2015 and 2016 data. The τ_h misidentification probability is shown as a function of the τ_h identification efficiency for six predefined working points (WPs), and compared to three WPs for the cone-based $\Delta\beta$ corrected isolation. [175]

A diverse set of simulated genuine-τ_h events (signal) from both SM and BSM processes with τ_h p_T ranging from 20 GeV to 2 TeV as well as simulated events with no genuine τ_h candidates present (background) are used for training the BDT algorithm. After training, τ_h identification and isolation is performed by simply placing a cut on the BDT discriminator.

The performance of the BDT discriminator is illustrated in Figure 8.7, where the τ_h misidentification probability is shown as a function of the τ_h identification efficiency. The six different data points correspond to predefined working points with different efficiencies and misreconstruction rates. Three working points for the cone-based $\Delta\beta$ corrected isolation are also shown for comparison.

In addition to jets misidentified as τ_h, efficiently reduced by the BDT discriminant, also electrons or muons can be misidentified as τ_h. Dedicated BDTs are also trained to identify and reject the candidates in both cases.

The anti-electron BDT uses a number of input variables describing the energy distribution in calorimeters (such as the share of energy between ECAL and HCAL as well as

the distance of strip photons from the leading track), the amount of bremsstrahlung emitted by the leading charged particle (based on both tracker and ECAL information) and the quality of the tracks when fitted with the standard method and with the GSF algorithm (optimized for electrons, see Section 8.1.5), among others. Also the basic properties of the τ_h candidate such as the mass, p_T and η are included. The details of the input variables and the training are given in Ref. [175].

The anti-muon BDT is based on inputs such as the number of hit and track segments in muon chambers geometrically matched with the τ_h candidate and the amount of calorimeter energy deposits and their distribution in the ECAL and HCAL. Also for this discriminator the details are given in Ref. [175].

8.2.6 Missing transverse momentum

The missing transverse momentum, $\vec{p}_\text{T}^{\text{miss}}$, is defined as the negative vector sum of the $\vec{p}_\text{T}$ of all reconstructed PF candidates [165]. This $\vec{p}_\text{T}^{\text{miss}}$ definition incorporates all the physics objects (muons, electrons, photons, τ_h candidates and jets), reconstructed from the PF candidates and identified as described in the previous sections, but also the unclustered energy, defined as the energy of all the PF candidates not clustered into any physics object. Therefore the estimation of $\vec{p}_\text{T}^{\text{miss}}$ is affected by spatial and energy resolutions of all PF candidates, and in addition to genuine momentum imbalance, it can be altered by mismeasurement and detector artifacts.

The $\vec{p}_\text{T}^{\text{miss}}$ is calibrated by propagating the effect of the jet energy corrections described in Section 8.2.3 into it. Corrections to the energy scales of other physics objects are negligible compared to the jet energy corrections, so they are ignored. The *Type-I corrected* $\vec{p}_\text{T}^{\text{miss}}$ is defined as:

$$\vec{p}_\text{T}^{\text{miss}} = \vec{p}_\text{T}^{\text{miss,uncorrected}} - \sum_\text{jets}(\vec{p}_\text{T}^{\text{corrected}} - \vec{p}_\text{T}^{\text{uncorrected}}), \qquad (8.4)$$

where $\vec{p}_\text{T}^{\text{corrected}}$ ($\vec{p}_\text{T}^{\text{uncorrected}}$) is a jet p_T before (after) the jet energy corrections. To suppress the effect of pileup jets, only the jets with the corrected p_T above 15 GeV are included in the sum. As the correction is designed for quark and gluon jets, the jets corresponding to electromagnetic showers from electrons and photons are removed by excluding jets that have $> 90\%$ of the jet energy deposited in the ECAL. For the same reason, jets containing global muons or standalone muons are excluded from the sum.

Chapter 9

Statistical methods

The statistical analysis aims to draw conclusions about the observed data, under some theoretical hypothesis. If there is no significant excess corresponding to the presence of a signal, how large signals can be excluded based on the observed data? Or if there is an excess in the data, how likely is it to originate from the signal modeled by a given signal model?

In both cases, the *parameter of interest* in the analysis is the amount of signal, represented by the *signal strength modifier μ*. The signal strength modifier μ is defined as a parameter that varies the signal yield, thus representing different signal hypotheses. If s (b) is the expected event yield for signal (background) events, the expected total yield is $\mu s + b$.

The standard workflow of hypothesis testing is followed: First, the null hypothesis is defined and a suitable *test statistic* is constructed. The observed value of the test statistic is calculated from data, and conclusions are drawn by comparing it to the expected distribution of the test statistic. In this chapter, these steps of the hypothesis testing workflow and the relevant concepts are presented, as applied in the analysis presented in this thesis.

In traditional "cut-and-count" experiments the test statistic was defined simply based on the expected and observed event yields, obtained after online and offline selections. Modern computing techniques enable a more powerful "shape analysis" approach, where a *summary statistic* is calculated from the selected events and used to derive the test statistic. The summary statistic can be any distribution that discriminates between the background and signal events, such as a reconstructed mass distribution or output of an MVA classifier. In the analysis presented in this thesis, the transverse mass as defined in Chapter 10 is used. We refer to these distributions of the summary statistic as *templates*.

In a shape analysis, the test statistic incorporates both the expected event yield $\mu s_i + b_i$ and the observed yield n_i in each bin of the summary statistic. The normalization of the signal in the templates (s_i) can be based on a specific theoretical model, or it can be arbitrary, as it is only an initial value for the fit to data. If the production cross section (σ) and the branching fraction to the final state in case ($\mathcal{B}$) are known from the SM (or from a specific BSM model), the signal can be normalized accordingly and the signal strength modifier represents deviation from the theory expectation. In case of more generic searches, such as the $H^{\pm}$ search presented in this thesis, it is more convenient to normalize the signal templates to an arbitrary initial value, such as $\sigma = 1\,\mathrm{pb}$, $\mathcal{B} = 100\%$.

In the frequentist paradigm, a probability is defined in an objective way as the relative frequency of an event in the limit of a large number of trials. On the other hand, in the Bayesian paradigm a probability is interpreted as a subjective degree of belief, thus changing as new information is obtained. In the statistical interpretation of most high-energy physics experiments, including this analysis, the frequentist approach is followed. The reason is that we aim to interpret each individual search result in an objective way and independent of previous results. In Bayesian approaches the conclusions would be dependent on prior probability distributions affected by subjective judgments and previous results.

9.1 Exclusion of signal

When no clear excess is observed in data compared to the background expectation, the goal of the statistical analysis is to set a limit on μ, which can be interpreted as a limit on $\sigma\mathcal{B}$ by taking into account the initial normalization of the signal template. Here we present a method for signal exclusion, which is agreed upon and used by the ATLAS and CMS experiments and documented in Ref. [177].

The probability to observe n_i events in a template bin i when the expected yield is $\mu s_i + b_i$ is given by the Poisson probability distribution. The combined probability for all bins is

$$\mathcal{L}(\{n_i\}|\mu) = \prod_i \frac{(\mu s_i + b_i)^{n_i}}{n_i!} e^{-(\mu s_i + b_i)}. \tag{9.1}$$

Here $\{n_i\}$ denotes to the ensemble of bins, with n_i events in bin i. This value of μ that maximizes this *likelihood function* for the observed data n_i is the maximum-likelihood value μ_{ML}.

9.1.1 Test statistic

To construct the test statistic, the likelihood function is normalized by its maximum-likelihood value $\mathcal{L}(\{n_i\}|\mu_{\mathrm{ML}})$. For computational effectiveness, maximization of this function can be performed as the minimization of the corresponding *log-likelihood* function. According to this *modified frequentist* criterion [177], the test statistic $\tilde{q}_\mu$ is defined as a *likelihood ratio*

$$\tilde{q}_\mu = -2\ln\frac{\mathcal{L}(\{n_i\}|\mu)}{\mathcal{L}(\{n_i\}|\mu_{\mathrm{ML}})}. \tag{9.2}$$

The minimization of $\tilde{q}_\mu$ is subject to two constraints: Firstly, $\mu \geq 0$ (and $\mu_{\mathrm{ML}} > 0$), reflecting the fact that the presence of a signal should always be manifested as additional events. Secondly, $\mu_{\mathrm{ML}} \leq \mu$, meaning that cases where the observed value μ_{ML} is smaller than signal hypothesis μ are considered as evidence against the signal hypothesis, whereas the cases where $\mu_{\mathrm{ML}} > \mu$ are still considered as supporting the signal hypothesis. In other words, the goal is to test whether the signal is *at least* as large as suggested by the hypothesis.

The expected distributions of $\tilde{q}_\mu$ can be produced by generating pseudoexperiments, sampling random event yields from the Poisson distribution with a mean of $\mu s_i + b_i$. The pseudoexperiments are performed under two assumptions, namely under the signal hypothesis ($\mu = \mu_{\mathrm{hyp}}$) and under the background-only hypothesis ($\mu = 0$), while the value of $\tilde{q}_\mu$ is in both cases calculated assuming $\mu = \mu_{\mathrm{hyp}}$. When normalized to unity, the distributions become *probability density functions* (pdfs) $f(\tilde{q}_\mu|\mu = \mu_{\mathrm{hyp}})$ and $f(\tilde{q}_\mu|\mu = 0)$ that can be used to extract p-values.

9.1.2 P-values

The *p-value* for the observed data represents the probability of finding data of equal or greater incompatibility with the null hypothesis, under the assumption that the null hypothesis is true. Therefore if the p-value is smaller than a predefined threshold α, the null hypothesis can be considered as excluded at a confidence level (C.L.) of $1 - \alpha$.

When excluding signal, the null hypothesis is that a signal with a strength modifier μ_{hyp} is present in data, and the p-value is calculated as

$$p_\mu = P(\tilde{q}_\mu \geq \tilde{q}_\mu^{\text{obs}} | \mu = \mu_{\text{hyp}}) = \int_{\tilde{q}_\mu^{\text{obs}}}^{\infty} f(\tilde{q}_\mu | \mu = \mu_{\text{hyp}}) d\tilde{q}_\mu. \tag{9.3}$$

The value of p_μ gives the probability P that one obtains a test statistic ($\tilde{q}_\mu$) at least as large as the observed value $\tilde{q}_\mu^{\text{obs}}$, even though a signal of strength μ (or larger) is present in data.

Alternatively, one can choose the background-only hypothesis as the null hypothesis, and calculate a p-value

$$p_b = P(\tilde{q}_\mu < \tilde{q}_\mu^{\text{obs}} | \mu = 0) = \int_{-\infty}^{\tilde{q}_\mu^{\text{obs}}} f(\tilde{q}_\mu | \mu = 0) d\tilde{q}_\mu \tag{9.4}$$

that gives the probability to obtain a test statistic $\tilde{q}_\mu$ at least as small as $\tilde{q}_{\text{obs}}$ due to statistical upward fluctuations in data, even though no signal is present.

9.1.3 CL$_S$ method and observed limit

In the modified frequentist method, also known as the CL$_s$ method [178, 179], instead of using p$_\mu$ itself for final judgment, a new p-value-like variable is defined as

$$\text{CL}_s = \frac{p_\mu}{1 - p_b}. \tag{9.5}$$

If a given signal hypothesis μ_{hyp} yields CL$_s = \alpha$, it is stated that signals larger than μ_{hyp} are excluded at $(1 - \alpha)$ C.L.. Typically, 95% C.L. (i.e. CL$_s = \alpha = 0.05$) is quoted to summarize the results, meaning that if the signal hypothesis were true, there would be only 5% probability of obtaining the observed limit (or an even lower one).

In practice, in the limit calculation the value of μ in Eq. (9.2) is adjusted until a value $\mu_{95\%}$ yielding CL$_s = 0.05$ is found, and $\mu_{95\%}\sigma\mathcal{B}$ is quoted as the observed 95% C.L. exclusion limit for $\sigma\mathcal{B}$. The CL$_s$ method is known to give conservative exclusion limits, so the true confidence level is actually even higher [179].

There are several reasons for adopting CL$_s$ instead of p_μ to for signal exclusion. For example, when p_μ is used, increasing the background expectation leads to *improved* exclusion of the signal hypothesis, whereas when CL$_s$ is used, it leads to *weakened*
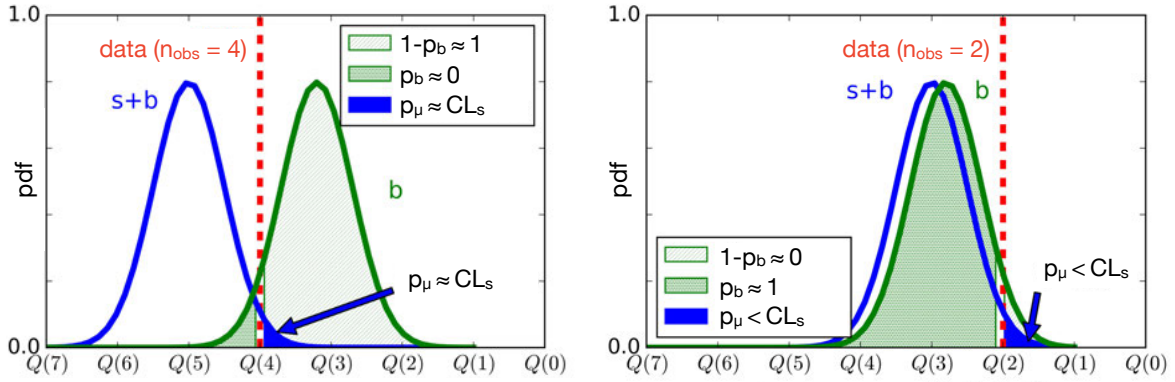
Figure 9.1: Illustration of CL_s and the probabilities p_μ and p_b in a case of a simple cut-and-count experiment, with signal hypotheses predicting a large signal yield (left) and a small yield (right). The probability density functions $f(\tilde{q}_\mu | \mu = \mu_{hyp})$ ("s+b") and $f(\tilde{q}_\mu | \mu = 0)$ ("b") are shown as a function of the test statistic value. $Q(n_{obs})$ denotes the value of the test statistic $\tilde{q}_\mu$ when n_{obs} events are observed in data. The dashed vertical line represents the event yield observed in data. [180]

exclusion. Similarly, with p_μ a mere downward fluctuation in data can lead to a strong exclusion limit, while CL_s is protected against the effect of such fluctuations. These and other arguments for the CL_s method are discussed in Ref. [179].

The relation between CL_s, p_μ and p_b is illustrated in Figure 9.1, showing the probability distributions of the test statistic $\tilde{q}_\mu$ in a case of large signal (left) and in a case of a small signal (right). In this example, a simple cut-and-count experiment with only one bin is considered, and $Q(n_{obs})$ denotes the value of the test statistic $\tilde{q}_\mu$ when n_{obs} events are observed in data. On the left, the data (shown as a red dashed line) are in good agreement with the background-only hypothesis, and agree poorly with the signal+background hypothesis. As a result, $1 - p_b \approx 1$, so $CL_s \approx p_\mu$, and since CL_s (corresponding to the blue area) is very small, the signal is excluded at high C.L.. On the right, the observed event yield (red line) is quite incompatible with both the background-only hypothesis and the signal+background hypothesis. While the value of p_μ is similar to the first case, the value of CL_s is much larger, so the signal is excluded at a considerably lower C.L. than in the first case.

9.1.4 Expected limits

Once the observed exclusion limit is calculated, it is natural to ask whether the result corresponds to what we would expect if no signal is present. For this purpose, a set of pseudoexperiments are performed and the exclusion limit is calculated for each of

them, following the method described above. These pseudoexperiments are created from Poisson distributions with means n_i, i.e. corresponding to the observed yields of data. The results are ordered by $\mu_{1-\alpha}$ and normalized to unity, producing a pdf.

When characterizing analysis sensitivity, typically *median* values, such as the median expected significance or median expected upper limits are quoted. The reason is that when a median value is propagated through any nonlinear (but monotonous) equation for another quantity, the result gives the median value for this other quantity, while for mean values this is only true with linear equations. The median expected limit corresponds to cumulative probability of 50% (half of the pseudoexperiments yield a smaller limit, half a larger limit). The 1σ band (also denoted as 68% expected band) corresponds to values of cumulative probability from 16 to 84%, and the 2σ band (95% expected band) corresponds to values from 2.5 to 97.5%.

9.1.5 Asimov data set and the asymptotic approximation

The expected limits can be used to estimate the sensitivity of the analysis already when it is in the blinded stage. For this purpose, the observed data is replaced by the *Asimov data set*, formally defined such that when it is used to evaluate the estimator for a parameter, one obtains the true parameter value [181]. The idea of the Asimov data set is to replace the ensemble of simulated (pseudoexperiment) data sets by a single representative one, so it is named after Isaac Asimov, the author of the short story "Franchise", where elections are held by selecting a single person to represent all voters.

The Asimov data set is a pseudo-data set where the data values n_i are set equal to the expectation $\mu_{\text{hyp}} s_i + b_i$ (in case of the background-only hypothesis, it is simply b_i).

As generating a large number of pseudoexperiments can be computationally heavy, a faster, approximate method is often used to obtain the distribution for the chosen test statistic. This *asymptotic approximation* is based on the observation that in the limit of large data samples, the profile likelihood ratio such as the Eq. (9.2) follows a noncentral chi-square distribution [181]. As a result, CL_{s} can be conveniently calculated as

$$\text{CL}_{\text{s}} \approx \frac{1 - \Phi(\sqrt{\tilde{q}_\mu^{\text{obs}}})}{\Phi(\sqrt{\tilde{q}_\mu^{\text{Asimov}}} - \sqrt{\tilde{q}_\mu^{\text{obs}}})}, \tag{9.6}$$

where $\tilde{q}_\mu^{\text{Asimov}}$ is the value of $\tilde{q}_\mu$ obtained using the background-only Asimov data set and $\Phi(x)$ is the cumulative Gaussian distribution

$$\Phi(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-x'^2/2} dx'. \tag{9.7}$$

In practice, the asymptotic approximation is found to produce accurate results even for relatively small data sets. In case of doubt, it is customary to validate the usage of this approximation by comparing the results with those obtained from the pseudoexperiments for a subset signal hypotheses.

9.2 Testing for a possible signal

In the case of an excess in data, its statistical significance needs to be quantified. Now the null hypothesis is the background-only hypothesis, so a suitable test statistic is obtained from Eq.(9.2) by setting $\mu = 0$:

$$\tilde{q}_0 = -2\ln \frac{\mathcal{L}(\{n_i\}|\mu = 0)}{\mathcal{L}(\{n_i\}|\mu_{\text{ML}})}. \tag{9.8}$$

The constraint $\mu_{\text{ML}} \leq \mu$ is required as previously.

The observed $\tilde{q}_0^{\text{obs}}$ is compared to the expected distribution by generating pseudoexperiments, calculating the pdf $f(\tilde{q}_0|\mu = 0)$ and extracting the local p-value as

$$p_0 = P(\tilde{q}_0 \geq \tilde{q}_0^{\text{obs}}|\mu = 0) = \int_{\tilde{q}_0^{\text{obs}}}^\infty f(\tilde{q}_0|\mu = 0) d\tilde{q}_0. \tag{9.9}$$

This p-value tells how likely it is to obtain $\tilde{q}_0$ at least as large as $\tilde{q}_0^{\text{obs}}$ via mere statistical fluctuations, without the presence of a signal. Using the asymptotic approximation, the value of p_0 can be estimated without any pseudoexperiments directly from the observed $\tilde{q}_0$ as $p_0 \approx 1 - \Phi(\sqrt{\tilde{q}_0})$.

The result of the test is often expressed as a significance Z, based on one-sided confidence interval of the Gaussian distribution $p_0 = \Phi(Z)$. Therefore in the asymptotic approximation, the observed significance is simply $Z \approx \sqrt{\tilde{q}_0}$.

For example, the 5σ significance (5 standard deviations in the Gaussian distribution) often quoted as a criterion for discovery of a new process, corresponds to the p-value of $p_0 = 2.87 \times 10^{-7}$.

9.3 Systematic uncertainties

The systematic uncertainties of the experimental setup and data processing methods enter the statistical analysis as a set of *nuisance parameters* $\vec{\theta}$. The nuisance parameters can alter signal and background yields in each bin of the templates, so when the uncertainties are taken into account, the event yields become functions of $\vec{\theta}$: $s_i \rightarrow s_i(\vec{\theta})$, $b_i \rightarrow b_i(\vec{\theta})$.

Every nuisance parameter θ has a *default value* θ_0 corresponding to the best estimate of the nuisance. In practice this corresponds to the nominal yield in the bin, not varied by the presence of the nuisance. Each nuisance parameter is assigned a probability density function $f(\theta|\theta_0)$, characterizing probabilities for variations around the default value. If the nuisance parameter affects all bins in a template in a similar way, multiplying the event yield by a given factor, we refer to it as *normalization uncertainty*. If the effect is defined separately for each template bin, so that the nuisance can affect both the shape and normalization of the distribution, it represents a *shape uncertainty*.

For the normalization uncertainties, a common choice for a probability density function is the log-normal pdf

$$f(\theta|\theta_0) = \frac{1}{\sqrt{2\pi}\ln\kappa}\exp(-\frac{(\ln(\theta/\theta_0))^2}{2\ln\kappa^2})\frac{1}{\theta}, \tag{9.10}$$

specified by the parameter κ. It affects the nominal event yield n such that the event yield modified by the nuisance has 68% probability to be within range $n/\kappa, n \times \kappa$. Therefore the modified event yield is always positive.

The shape uncertainties are estimated by varying the input parameter identified as the source of the uncertainty in the analysis and recalculating the results, resulting in shifted templates corresponding to $\pm 1\sigma$ variations of the input parameter. Then the pdf is constructed as a continuous set of variations by morphing the nominal and varied templates using a suitable morphing function. In some cases, simple linear interpolation can be sufficient, but usually a more sophisticated technique known as *horizontal morphing* [182] is used. This method is based on linear interpolation of the inverses of the cumulative distributions and has some elegant properties—for example, when two exponential decay distributions are interpolated with this method, the result corresponds to a linear interpolation of the decay constant.

Different sources of uncertainty, corresponding to different nuisance parameters, can be treated as fully correlated (100% correlation), anti-correlated (-100%), or independent (0%). The correct assignment of correlations depends on the specific uncertainties in hand. Partially correlated uncertainties are treated by splitting them to fully (un)correlated subcomponents.

The nuisance parameters enter Eq. (9.1) as follows, turning it into a *profile likelihood function* that depends both on the parameter of interest μ and on the full set of nuisance parameters $\vec{\theta}$:

$$\mathcal{L}(\{n_i\}|\mu, \vec{\theta}) = \prod_i \frac{(\mu s_i(\vec{\theta}) + b_i(\vec{\theta}))^{n_i}}{n_i!} e^{-(\mu s_i(\vec{\theta}) + b_i(\vec{\theta}))} f(\theta_0|\theta). \tag{9.11}$$

For a given μ, one can maximize the likelihood function to find the maximum-likelihood estimators $\vec{\theta}_{\mathrm{ML}}(\mu)$ for the nuisance parameters.

When these maximum-likelihood estimators are used to profile the nuisance parameters in the test statistic, Eq. (9.2) becomes

$$\tilde{q}_\mu = -2 \ln \frac{\mathcal{L}(\{n_i\}|\mu, \vec{\theta}_{\mathrm{ML}}(\mu))}{\mathcal{L}(\{n_i\}|\mu_{\mathrm{ML}}, \vec{\theta}_{\mathrm{ML}}(\mu_{\mathrm{ML}}))}. \tag{9.12}$$

When the limit is set with this test statistic, full information on the systematic and statistical uncertainties defined by the nuisance parameters is incorporated in the fit. For each nuisance parameter, deviation from the default value is allowed, but associated with a penalty in the fit as described by the pdf.

9.4 Statistical uncertainties

The limited number of events in each bin of the templates introduces statistical uncertainties that need to be taken into account. In the following, we discuss two different methods to achieve this: a traditional Barlow-Beeston method and an approximative *Barlow-Beeston-lite* approach [183] that has recently gained popularity in the CMS analyses.

In the traditional Barlow-Beeston method, each bin of each template is attached a separate nuisance parameter corresponding to the statistical uncertainty of the events in that particular bin. As the population of a bin follows a Poisson distribution, the uncertainty attached is simply the square root of the event count.

If the background consists of p processes with separate templates, and each template has n bins, np shape uncertainties are calculated by varying the event yield in each bin of each template up and down. As statistical fluctuations in bins are independent of each other, all these uncertainties are incorporated as np independent nuisance parameters in the likelihood.

As templates can have a large number of bins, this method often introduces a great number of shape uncertainties, and the limit calculation with full uncertainties can become computationally heavy. Therefore an approximative method, known as the Barlow-Beeston-lite approach, is often used instead. In this method, contributions from different background processes are summed into one total event count per bin, and this total count is assigned a single Gaussian-constrained uncertainty, reducing the number of new nuisance parameters to n.

In the Barlow-Beeston-lite approach the ML estimate of each nuisance parameter has an analytic formula that depends only on b_i (the total number of background events in template bin i) and n_i (the number observed events in bin i). Thus this part of the likelihood can be minimized analytically before proceeding to numerical minimization, which leads to reduced computation time and better fit stability.

PART III

Data Analysis

Chapter 10

Analysis strategy

In Chapters 10–13, the data analysis methods used to search for charged Higgs bosons decaying via $H^\pm \to \tau^\pm \nu_\tau$ are presented in detail. The results follow in Chapter 14. This search is the first CMS analysis on this decay channel based on 13 TeV collision data. A preliminary version of the analysis with a partial data set is documented in Ref. [72]. The final analysis and its results are published in Ref. [73].

As discussed in Section 3.2, the production of $H^\pm$ is possible via various processes and the mass of the $H^\pm$ determines the dominant process. In all production modes, the final state includes a $H^\pm$, a top quark, and possibly a b quark. The top quark decays further into a b quark and a W boson, which in turn can decay hadronically or leptonically.

In this analysis, we concentrate on the *hadronic final state* of the $H^\pm \to \tau^\pm \nu_\tau$ decay channel. In the hadronic final state, also denoted as the τ_h + jets final state, both the tau lepton (from $H^\pm$) and the $W^\pm$ boson (from a top quark) decay hadronically. As the total hadronic branching fraction is approximately 2/3 for both τ and $W^\pm$, almost half of all $H^\pm \to \tau^\pm \nu_\tau$ signal events evolve into the τ_h + jets final state. The leptonic final states where the tau lepton from the $H^\pm$ decay or the $W^\pm$ boson decay leptonically are targeted by a separate analysis and discussed briefly in Section 10.4.

Since the τ_h + jets final state looks similar regardless of the $H^\pm$ mass, the search can cover a wide $m_{H^\pm}$ range from 80 GeV up to 3 TeV. Examples of possible production processes, with the following decays that produce the τ_h + jets final state, are shown as tree-level Feynman diagrams in Figure 10.1. The first diagram shows the production of a light $H^\pm$ in a top quark decay, whereas the second shows the heavy $H^\pm$ production in association with a top quark (in the five-flavor scheme). In both cases, the final state contains the visible decay products of the hadronic tau lepton decay (τ_h), neutrinos

Figure 10.1: Tree-level Feynman diagrams for the production of a light $H^{\pm}$ in a top quark decay (left) and of a heavy $H^{\pm}$ in association with a top quark in the five-flavor scheme (right). In both cases, the decays producing a hadronic final state are shown. The $H^{\pm}$ decays to a tau lepton and a neutrino, and the tau lepton decays hadronically (τ_h). The top quark decays into a b quark (producing a b jet) and a $W^{\pm}$ boson that also decays hadronically (producing light quark jets) .

from the $H^{\pm}$ and τ decays, quarks from the W boson decay, and one or several b quarks.

The detector fingerprint of the fully hadronic final state can be easily inferred from these diagrams: regardless of the $H^{\pm}$ production mode, we expect the final state to contain a τ_h, missing transverse momentum (p_T^{miss}) due to neutrinos, and a number of jets including one or several b jets. No isolated electrons or muons are expected to be present in the targeted final state. The event selection criteria designed to extract this type of event topology are described in Chapter 11.

The physics objects are reconstructed using the methods described in Chapter 8, and used in the event selection. The exact identification criteria for different objects are described when the event selection is detailed in Chapter 11.

The special feature of the τ_h + jets final state is that most of the p_T^{miss} in the event originates from the neutrinos produced in the $H^{\pm}$ decay (and the subsequent tau lepton decay). Thus the τ_h candidate and the p_T^{miss} can be used to reconstruct the transverse mass (m_T) distribution of the tau-neutrino system, defined as

$$m_T(\vec{p}_T(\tau_h), \vec{p}_T^{\text{miss}}) = \sqrt{2p_T(\tau_h)p_T^{\text{miss}}(1 - \cos\Delta\phi(\vec{p}_T(\tau_h), \vec{p}_T^{\text{miss}}))}, \qquad (10.1)$$

where $\vec{p}_T(\tau_h)$ is the transverse momentum of the reconstructed τ_h and $\Delta\phi$ denotes a difference in azimuthal angle.

The selected *signal region* is contaminated by background events that contain a jet misidentified as a τ_h (jet $\rightarrow \tau_h$ events), most of which are QCD multijet events. This background contribution is estimated from a *control region* with a data-driven method presented in Chapter 12.

The top quark production processes in the SM ($t\bar{t}$ and single top quark production) also produce a significant *irreducible* background with a final state that is almost identical to the $H^{\pm}$ signal events. For example, changing the $H^{\pm}$ to a $W^{\pm}$ in Figure 10.1 describes one possible final state of $t\bar{t}$ production. These processes, as well as the SM processes with W or Z bosons (W+jets, Z/γ^*, WW, WZ, ZZ) produce events with genuine tau leptons that decay hadronically. The background from all these processes is referred to as the *genuine-tau background*. A data-driven method for estimating this background, known as *tau embedding*, is presented in Chapter 12. As this method is still considered experimental, it is not used for the final results, where instead the genuine-tau background is estimated from simulation. To suppress the genuine-tau background, the tau lepton helicity can be used to discriminate between tau leptons originating from $H^{\pm}$ and $W^{\pm}$ bosons, as discussed in Section 2.5.2.

Finally, a small background contribution arises from events with electrons or muons misidentified as τ_h (e/$\mu \rightarrow \tau_h$). This contribution is estimated from simulation. Also the background estimations from simulation are presented in Chapter 12.

10.1 Data

The analysis is based on the proton–proton collisions recorded by the CMS experiment during the year 2016. The center-of-mass energy of the collisions was 13 TeV. The amount of collision events delivered by the LHC in 2016 corresponds to an integrated luminosity of 41.07 fb^{-1}. The CMS detector successfully recorded 37.82 fb^{-1} (92%) out of it. Finally, 35.92 fb^{-1} (95%) of the recorded data were certified as high-quality data suitable for physics analyses, including this analysis. The pileup distribution in the data is shown in Figure 10.2. For the data collected in 2016, an average of approximately 23 interactions per bunch crossing was measured [109].

In 2015, the CMS detector recorded 3.80 fb^{-1} of data [109]. Since the additional statistical power obtained by including these data is negligible, while they would require separate calibration, optimization and background measurement efforts, the 2015 data are not included in this analysis.

Figure 10.2: The pileup distribution in the pp collision data recorded by the CMS experiment in 2016. [109]

Following the usual CMS procedure, in order to avoid any subconscious biases this analysis was designed, optimized, validated and reviewed by the Higgs Physics Analysis Group of the CMS Collaboration while keeping the data blinded in the signal region. Only after the analysis methods were fixed and initially approved, the data were unblinded and the final results as presented in Chapter 14 were extracted.

10.2 Signal modeling

The $H^\pm$ mass range from 80 GeV to 3 TeV is covered in the search, including the intermediate mass range near m_t.

The signal samples for the light $H^\pm$ mass values from 80 to 160 GeV are generated at next-to-leading order (NLO) with the MADGRAPH5_aMC@NLO v2.3.3 [137] generator, assuming $H^\pm$ production via top quark decay (pp $\rightarrow$ H$^\pm$W$^\mp$b$\overline{\text{b}}$). For the heavy $H^\pm$ mass range from 180 GeV to 3 TeV, the same approach is used, except that $H^\pm$ production via pp $\rightarrow$ tbH$^\pm$ is assumed, and the four-flavor scheme is used.

For the intermediate mass range from 145 to 200 GeV, the samples are generated at LO using the MADGRAPH5_aMC@NLO v2.3.3 with the model described in Ref. [70], which is available only at LO. The effect of using LO instead of NLO samples is estimated by comparing kinematic distributions and final event yields from LO and NLO samples where a specific production mode (either t $\rightarrow$ bH$^\pm$ or pp $\rightarrow$ tbH$^\pm$) is assumed. The observed difference in acceptance is corrected with LO-to-NLO

correction factors, described in detail later in Section 13.7. After the correction, the results from intermediate-range samples from 145 GeV up to 160 GeV are found to agree with those obtained using the light $H^\pm$ NLO samples, indicating that in this range other production modes than $t \to bH^\pm$ have a negligible effect on the acceptance. Similarly, in the 180–200 GeV mass range, the results agree with those obtained from the NLO samples assuming $pp \to tbH^\pm$. Therefore the NLO samples that assume a specific production mode are used for all mass points for which they are available, and the LO intermediate-range samples are used only for three mass points not covered by the NLO samples, namely 165, 170 and 175 GeV.

Two types of intermediate-mass signal samples are generated: so-called *no-neutral* samples where the contribution from $H^\pm$ production via a neutral Higgs boson resonance ($H \to H^\pm W^\mp$) are not taken into account, and *with-neutral* samples that account for these diagrams. In an ideal case, neutral Higgs loops should be taken into account in the signal prediction by using both types of samples with correct relative weights for each mass point. Unfortunately, the neutral Higgs loops prediction is always model dependent, since the loop contributions depend on the properties of neutral Higgs bosons that differ between different 2HDM scenarios. Thus only the model-independent no-neutral samples are used in the analysis. The kinematic distributions are found to be similar between the no-neutral and with-neutral samples, so omitting the with-neutral contribution does not considerably change the signal acceptance. This approach is also consistent with the NLO calculation for the reference cross sections, which are used to interpret the results. In this calculation the with-neutral contribution was also omitted, since the effect of the neutral Higgs boson contribution to the $H^\pm$ cross section is was found to be at the percent level [70].

For all signal samples up to $m_{H^\pm} = 500$ GeV, MADSPIN [184] is used to model the decay of the $H^\pm$, while PYTHIA 8.212 is used for the mass points above 500 GeV.

10.3 Simulation methods

The background events from the dominant $t\bar{t}$ production are generated using POWHEG v2.0 [138–140, 185, 186] with FxFx jet matching and merging [187]. The POWHEG v2.0 generator is used to model single top quark production via t-channel and tW production [188, 189], while the MADGRAPH5_aMC@NLO v2.2.2 generator is used for the

Table 10.1: Summary of the simulation software used for different processes.

Hard process	Generator
H$^{\pm}$	MADGRAPH5_aMC@NLO v2.3.3
t$\bar{\text{t}}$	POWHEG v2.0
Single top (t and tW channels)	POWHEG v2.0
Single top (s-channel)	MADGRAPH5_aMC@NLO v2.2.2
Z/γ^*	MADGRAPH5_aMC@NLO v2.2.2
W+jets	MADGRAPH5_aMC@NLO v2.2.2
Diboson (WW, WZ, ZZ)	PYTHIA 8.212
Subsequent steps (common for all samples)	Software
Parton showering and pileup simulation	PYTHIA 8.212
Detector simulation	GEANT4

s-channel production. The value of m_t is set to 172.5 GeV for all t$\bar{\text{t}}$ and single top quark samples.

The W+jets and Z/γ^* events are generated at LO using MADGRAPH5_aMC@NLO v2.2.2 with up to four noncollinear partons in the final state [190]. The diboson processes (WW, WZ, ZZ) are simulated using PYTHIA 8.212. The simulation of signal samples is detailed above in Section 10.2.

The generators are interfaced with PYTHIA 8.212 to model the parton showering, fragmentation, and the decay of the tau leptons. The PYTHIA parameters affecting the description of the underlying event are set to the CUETP8M1 tune [191] for all processes except t$\bar{\text{t}}$, for which a customized CUETP8M2T4 tune [192] is used. For all simulated signal and background samples, the NNPDF3.0 parton distribution functions [193] are used.

Generated events are processed through a simulation of the CMS detector based on the GEANT4 v9.4 software [152, 153], and they are reconstructed following the same algorithms as used for data. The effect of additional soft inelastic proton-proton (pp) interactions (pileup) is modeled by generating minimum bias collision events with PYTHIA and mixing them with the simulated hard scattering events. The effects from multiple inelastic pp collisions occurring per bunch crossing (in-time pileup), as well as the effect of inelastic collisions happening in the preceding and subsequent bunch crossings (out-of-time pileup) are taken into account.

A summary of the software used for different processes is shown in Table 10.1. The simulated events are weighted such that the final pileup distribution matches the one observed in data. Also several other corrections are applied to the simulated samples as described in Chapter 13.

10.4 Combination with the leptonic final state

The final states of all $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ signal events do not correspond to the τ_h + jets final state. Electrons or muons can be produced either from the decays of the tau leptons from $H^{\pm}$ decays, or from a $W^{\pm}$ boson decay. Thus the *leptonic final states* of the $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ process contain a single isolated lepton (electron or muon), p_T^{miss} (due to neutrinos), and hadronic jets. For maximal signal sensitivity, the results of the analysis presented in this thesis, targeting the hadronic final state, are combined with those from a separate analysis targeting the leptonic final states, detailed in Ref. [73].

In the analysis of leptonic final states, single electron and single muon triggers are used in the online selection. In the offline selection, events with a single isolated electron or muon and one, two or three hadronic jets are selected.

At least one of the jets is required to be b-tagged. Leptonic final states containing a hadronically decaying tau lepton (τ_h), originating from the $H^{\pm}$ or from the $W^{\pm}$ decay, and those without a τ_h are considered in separate categories. The latter case corresponds to events where the $H^{\pm}$ decays leptonically while the $W^{\pm}$ boson decays hadronically, or where either $H^{\pm}$ or $W^{\pm}$ decay produces a τ_h but it does not pass the identification criteria.

In the leptonic final states, the dominant background is $t\bar{t}$ production. This and other backgrounds are estimated from simulation. The number of QCD multijet events with jets misidentified as leptons is reduced to a negligible level by requiring a high p_T^{miss} and by applying angular selections based on $\Delta\phi(\ell, \vec{p}_T^{\text{miss}})$, $\Delta\phi(\text{leading jet}, \vec{p}_T^{\text{miss}})$, and $\min(\Delta\phi(\ell, \text{jet}_n))$, where jet_n refers to any of the 2–3 selected jets in the events.

The selected events are classified into several categories for statistical analysis, based on the presence or absence of a τ_h candidate, the jet multiplicity, the number of b-tagged jets and on the magnitude of p_T^{miss}. Together with the separate electron and muon final states, this results in 34 different categories.

The analyses of the hadronic and leptonic final states complement each other, because the relatively high trigger thresholds limit the sensitivity of the hadronic final state

Table 10.2: A summary of the phase space regions used in different $H^\pm$ analyses.

	No isolated electrons or muons	1 isolated electron or muon
2 jets	—	$H^\pm \to \tau^\pm \nu_\tau$, leptonic final states
3 jets	$H^\pm \to \tau^\pm \nu_\tau$, τ_h + jets final state	
> 3 jets		$H^\pm \to tb$, leptonic final states

in the low-$m_{H^\pm}$ region (below m_t), making the leptonic final states experimentally the most sensitive ones for the $H^\pm$ signal. In the high-$m_{H^\pm}$ region (above m_t) the hadronic final state dominates the sensitivity, since the selection efficiency is higher as a result of more inclusive jet multiplicity requirements.

The object definitions and event selection requirements used in the hadronic and leptonic analyses are coordinated to ensure that the analyses use disjoint phase space regions, so that they have no events in common after all selections. This way the statistical combination of the results is straightforward as the data used in the two analyses are uncorrelated. In practice, this is ensured by common definitions in the lepton identification and isolation: the event selection in the hadronic channel rejects all events with isolated electrons and muons, whereas the leptonic analysis only uses these events. In the analysis of leptonic final states, events with more than 2–3 jets are not selected, because the high jet multiplicity events are expected to be more sensitive in the $H^\pm \to tb$ decay channel. In order to allow statistical combination of the results from $\tau^\pm \nu_\tau$ and tb channels, the jet selection criteria are used to ensure that the two analyses do not use the same events. The different phase space regions defined by the lepton isolation and the jet multiplicity, used by the different analyses, are summarized in Table 10.2.

Shared object definitions, such as identification algorithms, isolation criteria and working points, also mean that most sources of systematic uncertainties are shared between the hadronic and leptonic final states of the $H^\pm \to \tau^\pm \nu_\tau$ channel. Correlation of these systematic uncertainties in the combined statistical analysis allows the data from one final state to constrain the systematic uncertainties in another final state.

In Chapter 14, after the final results of the analysis of the τ_h + jets final state are presented, also the combined results from the hadronic and leptonic final states are shown and the contribution of each analysis in the final results is discussed.

Chapter 11

Event selection

This chapter describes the consecutive event selection steps applied in order to maximize the signal sensitivity. Table 11.1 summarizes the online and offline selection steps described in Sections 11.1 and 11.2, with the final selection thresholds optimized as described in Section 11.3. The resulting signal selection efficiency is reported in Section 11.4.

Table 11.1: A summary of the online and offline event selection criteria. Selection steps 3–5 are referred to as *baseline selections*. Steps 6–8 efficiently reduce the background from jets misidentified as τ_h, and step 9 discriminates between genuine taus from signal and backgrounds events.

	Selection	Description
1	$\tau_h + p_T^{miss}$ trigger	Selection of signal-like events using a dedicated trigger
2	Data quality filters	Events required to pass the p_T^{miss} filters
3	τ_h identification	$\geq 1\ \tau_h$ candidates with $p_T > 50\,\text{GeV}$, $|\eta| < 2.1$ (loose MVA τ_h ID)
4	Lepton veto	No isolated electrons (muons) with $p_T > 15(10)\,\text{GeV}$, $|\eta| < 2.5$
5	Jets selection	≥ 3 jets with $p_T > 30\,\text{GeV}$, $|\eta| < 4.7$ and passing the loose jet ID
6	b jet selection	≥ 1 b-tagged jet with $p_T > 30\,\text{GeV}$ and $|\eta| < 2.4$ (CSVv2 medium WP)
7	p_T^{miss}	Type-I corrected $p_T^{miss} > 90\,\text{GeV}$
8	Angular selection	$R_{bb}^{min} > 40°$ to reduce the jet $\to \tau_h$ background
9	R_τ categorization	Reconstruct m_T separately for events with $R_\tau > 0.75$ and $R_\tau < 0.75$

11.1 Online event selection

In the Level-1 trigger, events are selected by requiring the Level-1 p_T^{miss}, calculated as described in Section 6.1.4, to be larger than $80\,\text{GeV}$.

In the HLT, an algorithm requiring the presence of a τ_h candidate with a loose track-based isolation, and missing transverse momentum estimated from calorimeter information ($p_T^{\text{miss,calo}}$) is applied.

The HLT τ_h candidate is reconstructed using a fast cone-based algorithm that proceeds through several steps (levels), rejecting a fraction of the τ_h candidates at each level. First, at *level 2*, the τ_h candidates are identified using only the calorimeter energy deposits, and are required to pass a minimum energy threshold. At *level 2.5*, a loose isolation is required, based on tracks reconstructed from hits in the pixel detector and associated with the τ_h decay vertex. The τ_h candidates passing the isolation, as well as the candidates for which no vertex was reconstructed, are passed to the next level for additional scrutiny. Finally, at *level 3*, the full reconstruction of tracks and PF candidates is carried out, including the reconstruction of anti-k_T jets with distance parameter 0.4. The compatibility of each jet with the τ_h hypothesis is tested, similarly to the HPS algorithm used offline (Section 8.2.5) but with considerably looser selection criteria. The photons contained in a jet are clustered into $\eta \times \phi = 0.05 \times 0.02$ strips, which are combined with charged hadrons found inside a signal cone with $\Delta R = 3.6\,\text{GeV}/p_T$ (where ΔR is allowed to vary from 0.08 to 0.12).

As the final step, a cone-based isolation is applied by identifying the charged hadrons associated with the τ_h vertex but not belonging to the τ_h candidate within an isolation cone of $\Delta R < 0.4$. The p_T sum of these charged hadrons is required to be less than 3 GeV. To improve the selection efficiency, the isolation requirement is relaxed by 6% for τ_h candidates with $p_T > 100\,\text{GeV}$. This relaxation can be allowed without increasing the trigger rate too much, because the number of jets misidentified as τ_h passing the selections decreases with p_T. A detailed description of τ_h reconstruction and isolation in the HLT can be found in Ref. [175].

The HLT τ_h candidates passing all the above steps, with a reconstructed transverse momentum $p_T > 50\,\text{GeV}$ and with a leading track transverse momentum (from the leading charged hadron candidate) $p_T^{\text{track}} > 30\,\text{GeV}$ are selected. As the HLT algorithm uses tracker information, and because the tracker extends only up to $|\eta| = 2.5$ and an isolation cone of $\Delta R = 0.4$ is used, the τ_h candidates can be reconstructed only for $|\eta| < 2.1$.

For the p_T^{miss} part of the HLT algorithm, the $p_T^{\text{miss,calo}}$ is computed simply as the negative vector sum of the transverse energies of all calorimeter towers. It is required to be larger than 90 GeV.

With these trigger thresholds, the $\tau_h + p_T^{\text{miss}}$ signal trigger produced a rate of approximately 20 Hz during a typical LHC fill with an instantaneous luminosity of $1.4 \cdot 10^{34} \, \text{cm}^{-2} \, \text{s}^{-1}$ in 2016 [175].

11.1.1 Trigger efficiency

The efficiency of the τ_h part of the trigger is determined with the tag-and-probe technique [132], using $Z/\gamma^* \to \tau^+\tau^-$ events with one hadronic and one muonic tau lepton decay.

These events are selected with a single muon monitoring trigger, requiring a muon with $p_T > 21 \, \text{GeV}$ and $|\eta| < 2.1$. The tag-and-probe method is applied, using an isolated muon selected with the monitoring trigger ("tag") and a τ_h candidate ("probe") with $p_T > 20 \, \text{GeV}$, $|\eta| < 2.1$ and separated from the muon by $\Delta R(\tau_h, \mu) > 0.4$. In addition, at least two jets with $p_T > 30 \, \text{GeV}$ are required. The other identification criteria for objects are similar to those used in the signal selection (Section 11.2.2). The $Z/\gamma^* \to \mu\mu$ events are suppressed by requiring $20 < m_{\text{inv}}(\mu, \tau_h) < 80 \, \text{GeV}$ while the W+jets events are reduced with the requirement that $m_T(\mu, \vec{p}_T^{\text{miss}}) < 40 \, \text{GeV}$.

As the trigger efficiency is different for τ_h candidates from genuine taus and from misidentified jets or leptons, a high genuine-τ_h purity of the selected sample is important for an unbiased measurement. The purity of events with genuine taus with this selection is found to be $> 90\%$.

The measured HLT efficiency for the tau part of the trigger is shown in Figure 11.1 as a function of the τ_h p_T. The efficiency varies between 50 and 100%, as a function of p_T and η of the τ_h. As the efficiency measured from data differs from the efficiency obtained using simulated samples, the latter are corrected for the difference as detailed later in Section 13.1.1.

The efficiency of the $p_T^{\text{miss,calo}}$ part of the trigger is measured using a prescaled single-τ_h monitoring trigger, identical to the τ_h part of the signal trigger but with no $p_T^{\text{miss,calo}}$ requirement. The efficiency is defined as the number of the events passing both the signal trigger (with both τ_h and $p_T^{\text{miss,calo}}$ selection requirements) and the monitoring trigger, compared to the number of events passing the monitoring trigger (with only the τ_h selection requirements). To select events with a signal-like topology, all baseline selection steps as described in Section 11.2.2 are applied in the offline selection.

Figure 11.1: Trigger efficiency for the τ_h part of the trigger, as a function of the offline recon-structed p_T of the τ_h candidate, for data and simulated events. The corresponding fit functions and the ratio of data and simulation are shown. The offline selection requirement of $p_T > 50$ GeV is illustrated with a light grey area on the left.

The L1+HLT efficiency of the p_T^{miss} part of the trigger is shown in Figure 11.2. and varies between 10 and 100%, depending on the value of the p_T^{miss}. Similarly to the τ_h trigger case, a correction for simulated events is derived as detailed in Section 13.1.1.

To reduce uncertainties in the efficiencies from limited event yields, the measured efficiencies are fitted using suitable fit functions. This way the information gained in the phase regions with large event yields can be exploited to constrain the efficiency estimates in less populated regions. Binned maximum likelihood fits are performed for data and simulation separately. For data (simulation), the Sigmoid function (Crystal ball cumulative distribution function) is used for the τ_h part of the trigger, while the Richards function (Sigmoid function) is used for the p_T^{miss} part. The method was developed in Ref. [194], where also the definitions of the fit functions can be found. The fit results are shown in Figures 11.1 and 11.2.

Figure 11.2: Trigger efficiency for the p_T^{miss} part of the trigger, as a function of offline reconstructed p_T^{miss} (with Type 1 corrections applied), for data and simulated events. The corresponding fit functions and the ratio of data and simulation are shown. The offline selection requirement of $p_T^{miss} > 90$ GeV is illustrated with a light grey area on the left.

11.2 Offline event selection

11.2.1 Data quality filters

As the identification of the τ_h + jets final state relies heavily on the correct reconstruction of p_T^{miss}, the first step in the offline event selection is to filter out events that are likely to contain anomalous p_T^{miss}, arising from reconstruction failures or detector malfunction rather than the true particle content of the event.

For this purpose, a set of data quality filters is applied. They include calorimeter filters that identify and reject events with calorimeter anomalies affecting the momentum balance. These anomalies include abnormal noise in the HCAL, exceptionally high energies in the ECAL superclusters, or temporarily nonfunctional calorimeter towers.

Separate filters are used to identify events affected by a *beam halo* from muons produced in collisions upstream of the detector, causing characteristic energy deposits in calorimeters and CSC muon stations. A third set of filters suppresses misreconstruction effects by looking for high-p_T tracks with low track quality, which often correspond to misreconstructed muons and charged hadrons. Combined, these filters are typically able to reject 90% of the spurious-p_T^miss events, with per mille level mistagging rate. The data quality filters are described in detail in Ref. [165].

11.2.2 Baseline selections

In the offline selection, low thresholds for the p_T of the reconstructed τ_h and p_T^miss are needed to maximize the sensitivity for light $\text{H}^\pm$ (with $80 < m_{\text{H}^\pm} < 160\,\text{GeV}$). Thus selection criteria identical to those in the HLT are applied to the reconstructed τ_h candidate and to the offline-reconstructed (PF) Type-I corrected p_T^miss.

The fact that identical selection thresholds are used both in trigger and in offline selection for τ_h p_T and p_T^miss means that the events in the turn-on part of the trigger efficiency enter the offline analysis. Therefore the correct estimation of the turn-on is important for the correct normalization of simulated events.

The one-prong τ_h candidates, corresponding to τ decays into a charged pion and up to two neutral pions, are selected for further analysis. The three-prong τ_h candidates are rejected, as they are found to bring only a small improvement in sensitivity, while they would require a separate jet $\to \tau_\text{h}$ background estimation as the fake factors are different between 1-prong and 3-prong τ_h candidates.

At least τ_h candidate passing a loose working point of the τ_h identification MVA discriminant (Section 8.2.5) is required. The selected working point corresponds to an overall τ_h identification efficiency of $\approx 50\%$ (determined from $\text{Z}/\gamma^* \to \tau^+\tau^-$ events) and 3×10^{-3} probability for misidentifying a jet as a τ_h (determined from QCD multijet events).

Events are required to contain at least three jets with $p_\text{T} > 30\,\text{GeV}$ and $|\eta| < 4.7$, separated from the reconstructed τ_h by $\Delta R > 0.5$. Loose jet identification criteria as described in Ref. [167] are applied.

At least one of the jets is required to pass the b jet identification performed with the CSVv2 algorithm described in Section 8.2.4, with $|\eta| < 2.4$. The working point for the multivariate discriminant is chosen such that the probability to misidentify jets

originating from light-flavor quarks or gluons as b jets is 1%, corresponding to 65% efficiency for the selection of genuine b jets.

As no isolated electrons or muons should be present in the fully hadronic final state, any event with isolated electrons (muons) with $p_T > 15(10)\,\text{GeV}$ and $|\eta| < 2.5$ is rejected. For muons, loose identification criteria as described in Section 8.2.1 are applied. For electrons, a loose working point with a 95% efficiency is used in the MVA identification described in Section 8.2.2. For both muons and electrons, the isolation criterion is defined by requiring the p_T sum of the PF candidates within the mini-isolation cone (as defined in Section 8.2.1) to be less than 40% of the electron/muon p_T. The presence of leptonic $W^\pm$ decays from the top quark would lead to a smearing of the high edge of the m_T distribution for the $t\bar{t}$ and W+jets backgrounds. The lepton veto also ensures that there is no overlap in the selected events with respect to the analysis targeting the leptonic final state, allowing a statistical combination of the results from the hadronic and leptonic final states.

All these selection steps, excluding the offline p_T^{miss} and b jet selections, are collectively referred to as *baseline selections*. The baseline selections loosely ensures that the collection of physics objects in the selected events corresponds to the detector fingerprint expected from the signal events, but does not yet suppress the dominant jet $\to \tau_h$ background. After the baseline selections, the b jet selection efficiently reduces both jet $\to \tau_h$ and W+jets backgrounds, and the p_T^{miss} selection further suppresses the jet $\to \tau_h$ events.

Finally, as described in the following, angular selections are applied to reduce the dominant jet $\to \tau_h$ background, while the categorization based on the R_τ variable is used to discriminate between the genuine-tau background events (mostly from $t\bar{t}$ production) and the $H^\pm$ signal events.

11.2.3 Angular selection

After the baseline selections, the jet $\to \tau_h$ background is dominated by QCD multijet events where a jet misidentified as a τ_h is in a back-to-back configuration with the p_T^{miss} arising from an incorrect estimation of the jet momenta.

Thus the amount of QCD multijet events can be reduced by placing a selection on the $\Delta\phi$ difference between the τ_h and p_T^{miss}. But the definition of m_T (Eq. (10.1)) implies that the back-to-back configuration ($\Delta\phi = 180°$) corresponds to large values of m_T,

where also most of the signal is expected, especially for large $H^\pm$ mass values. Thus a simple $\Delta\phi$ cut would also suppress the signal significantly.

The problem is solved by taking also into account the jet directions. Typically QCD multijet events that pass the p_T^{miss} cut and enter the signal region are dijet events with two high-p_T back-to-back jets. If the momentum of one jet is overestimated and the other one is underestimated, the resulting momentum imbalance is interpreted as large p_T^{miss}. The defining feature of these QCD multijet events is that the p_T^{miss} is collinear in ϕ with one of the leading jets, which is rare for events with p_T^{miss} arising from neutrinos, such as $H^\pm$ signal or $t\bar{t}$ events.

Thus the jet $\rightarrow \tau_h$ background can be efficiently suppressed with a discriminant that requires the τ_h and p_T^{miss} to be back-to-back ($\Delta\phi(\tau_h, \vec{p}_T^{miss})$) and one of the leading jets to align with p_T^{miss} ($\Delta\phi(\text{jet}, \vec{p}_T^{miss})$). This discriminant is defined in the ($\Delta\phi(\tau_h, \vec{p}_T^{miss})$, $\Delta\phi(\text{jet}, \vec{p}_T^{miss})$) plane as

$$R_{bb}^{min} = \min_n \left\{ \sqrt{\left(180° - \Delta\phi(\tau_h, \vec{p}_T^{miss})\right)^2 + \left(\Delta\phi(\text{jet}_n, \vec{p}_T^{miss})\right)^2} \right\}, \qquad (11.1)$$

where the index n runs over the three highest p_T jets (jet_n) in the event.

The selected events are required to have $R_{bb}^{min} > 40°$. The effect of this cut on the ($\Delta\phi(\tau_h, \vec{p}_T^{miss})$, $\Delta\phi(\text{jet}, \vec{p}_T^{miss})$) plane and the regions where different types of events (signal events, irreducible $t\bar{t}$ background and the QCD multijet events) are mostly concentrated are schematically shown in Figure 11.3. The distribution of the R_{bb}^{min} variable after applying all other selections is shown in Figure 11.4.

11.2.4 Categorization of events

While the jet $\rightarrow \tau_h$ background can be efficiently suppressed with the b jet, p_T^{miss} and R_{bb}^{min} selections, and also the W+jets background is reduced with the b jet identification, the $t\bar{t}$ and single top production backgrounds with genuine tau leptons, b jets and p_T^{miss} are largely irreducible.

However, the τ helicity correlations can be utilized to suppress these backgrounds. As discussed in Section 2.5.2, an event sample enriched with $H^\pm$ signal events can be selected by imposing a cut on $R_\tau = p_{track}/p_{\tau_h}$, where p_{track} is the reconstructed three-momentum of the leading charged particle in the τ_h candidate, and the p_{τ_h} is the three-momentum of the τ_h candidate itself.
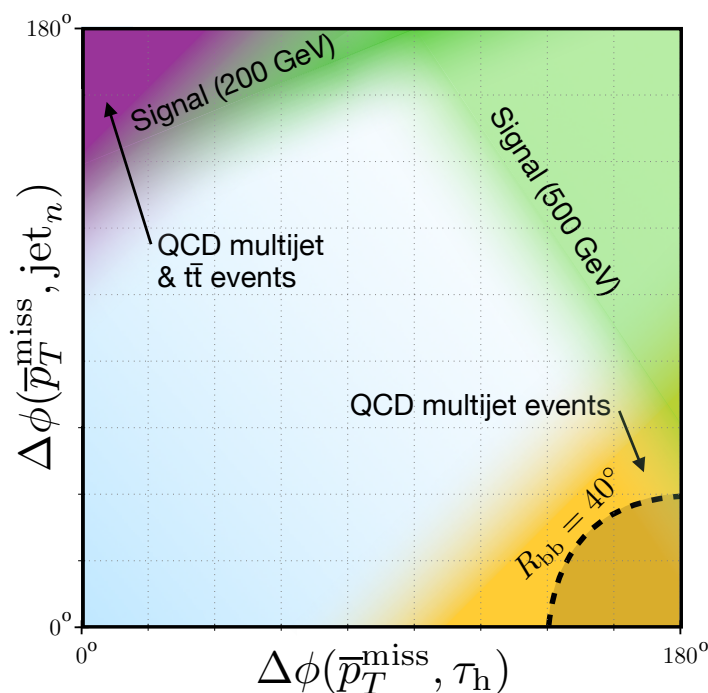
Figure 11.3: A schematic picture of the $(\Delta\phi(\tau_h, \vec{p}_T^{miss}), \Delta\phi(jet, \vec{p}_T^{miss}))$ plane used in the R_{bb}^{min} selection. The selection affects the bottom-right corner of the plane, containing mostly QCD multijet events. Also the regions populated with signal events (using $m_{H^\pm} = 200\,\text{GeV}$ and $500\,\text{GeV}$ as examples) and with the irreducible $t\bar{t}$ background are shown for illustration.

Therefore after all selections, the selected events are classified into two categories based on the value of the variable $R_\tau = p_{track}/p_{\tau_h}$. The distribution of the R_τ variable is shown in Figure 11.5. After all other selections, most of the signal events have a large value of R_τ as expected, and the high-R_τ category provides a good signal-to-background ratio. However, for TeV-range $m_{H^\pm}$ hypotheses, the signal events are more evenly distributed between the two categories, so the inclusion of the background-dominated low-R_τ category in the statistical analysis further improves the overall signal selection efficiency.

At the other end of the mass range, including the low-R_τ category also improves sensitivity for the light $H^\pm$ mass hypotheses of 80–90 GeV, as it allows the data to constrain the systematic uncertainties related to simulated backgrounds containing $W^\pm$ bosons, which is necessary to distinguish a signal situated almost on top of the $W^\pm$ peak in the m_T distributions.

Figure 11.4: The distribution of the angular discriminant R_{bb}^{min} after applying all other selections including the $R_\tau = p_{track}/p_{\tau_h} > 0.75$ requirement.

Separating the two categories at $R_\tau = 0.75$ maximizes the signal sensitivity across the $m_{H^\pm}$ range. In the category defined by the $R_\tau > 0.75$ cut, the leading charged particle is required to carry a large fraction (>75%) of the visible τ energy, while the other category is defined by the inverse requirement of $R_\tau < 0.75$.

11.3 Optimization of the event selection

The event selection was tested and optimized using the simulated signal and background samples described in Sections 10.2 and 10.3 and the data-driven estimation method for the jet $\to \tau_h$ background described in Section 12.3, while keeping the data blinded in the signal region. The selection thresholds used in the Run-1 analysis based on 8 TeV data (Ref. [27]) for the same channel were chosen as the initial cut values for optimization. To find the optimal selection criteria, the thresholds of kinematic cuts and working points of discriminators were varied in a systematic way. As the different

Figure 11.5: The distribution of the R_τ variable used for categorization after applying all other selections including the $R_{bb}^{min} > 40°$ requirement.

11.4 Signal selection efficiency

The signal selection efficiency obtained with the above offline and online selections is illustrated in Table 11.2, showing the estimated number of signal events after each selection step. The signal yields are shown for the mass hypotheses $m_{H^\pm}$ = 120, 170, 500, and 1000 GeV. In these yields, all corrections applied to simulated samples, detailed later in Section 13, are taken into account. The signal samples are normalized

Table 11.2: Number of selected events as a function of applied selections and including scale factors (SF) for signal mass hypotheses of $m_{H^\pm}$ = 120, 170, 500, and 1000 GeV. All signal samples are normalized to a production cross section of 1 pb, assuming $\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau) = 1$. Tau selection includes the $R_\tau > 0.75$ cut. The statistical uncertainties are also shown.

$m_{H^\pm}$ (GeV)	120	170	500	1000
All events	35919 ± 45	35919 ± 14	35919 ± 47	35919 ± 15
Trigger	845.7 ± 9.4	3402.1 ± 4.6	11310 ± 49	21064 ± 24
Tau selection	92.5 ± 3.1	509.0 ± 1.8	2334 ± 22	3367.9 ± 9.4
Trigger & tau id. SF	62.2 ± 2.1	369.9 ± 1.3	1947 ± 18	3045.0 ± 8.6
Lepton veto	42.5 ± 1.7	329.3 ± 1.2	1463 ± 16	2242.3 ± 7.3
≥ 3 hadronic jets	28.0 ± 1.4	218.4 ± 1.0	843 ± 13	1981.5 ± 7.0
≥ 1 b-tagged jets	21.6 ± 1.2	143.2 ± 0.8	623 ± 11	1274.9 ± 5.6
b-tagging SF	20.9 ± 1.2	140.4 ± 0.8	612 ± 11	1253.9 ± 5.5
$p_T^{miss} > 90$ GeV	18.2 ± 1.1	130.5 ± 0.8	599 ± 11	1245.4 ± 5.5
Angular selections	17.4 ± 1.1	124.6 ± 0.8	449.3 ± 9.3	906.4 ± 4.7
NLO/LO SF		81.0 ± 0.5		
Final event yield	17.4 ± 1.1	81.0 ± 0.5	449.3 ± 9.3	906.4 ± 4.7

to a production cross section of 1 pb, assuming $\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau) = 1$. The yields are tabulated for the $R_\tau > 0.75$ category. Also the statistical uncertainties due to the limited sample size are shown.

We observe that the selection efficiency strongly depends on the $H^\pm$ mass. For $m_{H^\pm} = 120$ GeV, the typical values for the p_T of the τ_h and the p_T^{miss} correspond to the trigger turn-on, so the majority of signal events is lost already in the online selection. The probability of the event to pass the offline selection requirements for τ_h p_T and p_T^{miss} increases with $m_{H^\pm}$.

A similar table showing the selection efficiencies for the simulated background processes will be presented at the end of Section 12, after the background estimation methods have been presented.

Chapter 12

Background estimation

Accurate estimation of the background from the Standard Model processes is an essential part of the analysis, as we aim to distinguish a possible small excess caused by signal events from this background. The dominant background processes in the selected τ_h + jets final state are the QCD multijet and $t\bar{t}$ production. Other backgrounds are single top quark production, W boson production in association with jets, Z/γ^* processes, and diboson production. We refer to $t\bar{t}$ and single top quark events as "top events", and to W+jets, Z/γ^*, and diboson events as "electroweak events".

The backgrounds can be classified into three types based on the origin of the τ_h candidate in the event: genuine-tau events, jet $\to \tau_h$ events, and e/$\mu \to \tau_h$ events. According to simulation studies, the majority of the selected top and electroweak events contain a τ_h candidate that corresponds to a genuine tau lepton decaying hadronically. The background from these events is referred to as the genuine-tau background.

In this analysis, the genuine-tau background is estimated from simulation as described in Section 12.1. An alternative method for estimating this background known as *tau embedding*, which is based on single-muon events and uses only minimal input from simulation, is presented in Section 12.2. An earlier version of the tau embedding method was developed in Ref. [195] and used in the CMS analysis of the 8 TeV data [27]. Herein, a new version of the method is presented. As the background estimate obtained with tau embedding relies mostly on data, it can describe accurately some features of the events that are challenging to simulate correctly, such as the production of associated jets and the underlying event. On the other hand, the power of the method is limited by its systematic uncertainties and by the number of single-muon

events available for embedding. As this method is still considered preliminary, it is not used for the final results presented in Chapter 14.

Another large source of τ_h candidates are events where a jet is misidentified as τ_h. These jet $\to \tau_h$ events are mostly QCD multijet events, but they also include a fraction of top and electroweak events. As the QCD multijet background cannot be estimated from simulation in a reliable way, it is estimated using a data-driven method. A *fake factor method* developed in Ref. [196] is applied to estimate this background as described in Section 12.3.

Finally, in rare cases also electrons or muons can be misidentified as τ_h. This e$/\mu \to \tau_h$ background is estimated from simulation as described in Section 12.4. The final background estimates are summarized in Section 12.5.

12.1 Estimation of genuine-tau background from simulation

The backgrounds from top and electroweak events containing a genuine τ_h are estimated from simulation, using the simulated background samples listed in Section 10.3. The methods used in the simulation are presented in Chapter 7. The correct identification or misidentification of τ_h is determined by requiring a generator-level tau lepton to match with the reconstructed τ_h within a ΔR cone of 0.1.

The simulated background samples are normalized to the theoretical cross sections for the corresponding processes. For the $t\bar{t}$ background and the single top quark background in the s and tW channels, the cross sections are calculated at next-to-NLO precision [197, 198]. NLO precision calculations are used for single top quark production in the t channel, and for the W+jets, Z$/\gamma^*$, and diboson processes [198–201].

For the dominant $t\bar{t}$ background, the estimate was cross-checked by comparing the results from two generators: POWHEG v2.0 and MADGRAPH5_aMC@NLO v2.2.2. The comparison is illustrated in Figure 12.1, where the transverse mass distributions obtained from $t\bar{t}$ event samples produced with these two generators and processed through the full reconstruction and analysis chain are shown for the category with $R_\tau > 0.75$. The two distributions are found to be compatible both in shape and in terms of the overall normalization. For final results, the simulated samples generated with POWHEG v2.0 are used, since they offer better statistical precision. Due to negative event weights arising from the NLO generation with MADGRAPH5_aMC@NLO, a much

Figure 12.1: Comparison of the $t\bar{t}$ background estimates obtained using two different generators: POWHEG v2.0 and MADGRAPH5_aMC@NLO v2.2.2. The transverse mass distributions from both samples are normalized according to the latest theory estimate at next-to-NLO precision. The two distributions are compatible within the statistical uncertainties.

larger amount of simulated events would be needed to obtain a similar precision with this generator. The $t\bar{t}$ background prediction from POWHEG has also been thoroughly compared to the distributions from data as part of the CMS measurements of the $t\bar{t}$ inclusive and differential cross sections [202, 203].

The various systematic uncertainties related to the genuine-tau background are discussed later in Chapter 13.

12.2 Estimation of genuine-tau background from data

In the Standard Model, the couplings of different leptons to the gauge bosons are flavor-independent. From this *lepton universality* it follows that $\mathcal{B}(W^{\pm} \to \tau^{\pm} \nu_{\tau}) = \mathcal{B}(W^{\pm} \to \mu^{\pm} \nu_{\mu})$ and $\mathcal{B}(Z \to \tau\tau) = \mathcal{B}(Z \to \mu\mu)$. As the tau leptons in the genuine-tau background events are produced either in $W^{\pm}$ decays ($t\bar{t}$, single top and W+jets, WW and WZ events) or in Z boson decays (Z/γ^*, WZ and ZZ events), it follows that the yield of the genuine-tau background can be estimated using single-muon events. As Higgs bosons couple to massive particles more strongly than to lighter particles,

$\mathcal{B}(H^{\pm} \to \mu^{\pm}\nu_{\mu},) \ll \mathcal{B}(H^{\pm} \to \tau^{\pm}\nu_{\tau})$, and the signal contamination from these events is negligible.

As the methods used in triggering, reconstruction, identification and isolation are different for muons and τ_h candidates, all these differences must be taken into account to obtain a correct estimate. The *tau embedding* method offers a straightforward way to achieve this. In tau embedding, a suitable data sample of single-muon events is selected. The energy deposits of reconstructed muons are removed from the event, and the muon is replaced with a simulated tau lepton, which has kinematic properties similar to the removed muon. This results in a sample of *hybrid events*, which are real data events except that they contain a simulated tau lepton.

A background estimate based on hybrid events can be considered more reliable than an estimate fully based on simulation, since the only input from simulation needed in the embedding method is the tau lepton. The tau lepton decay is a theoretically well understood process and can be described by simulation quite accurately. Other parts of the hybrid event which are more challenging to simulate accurately, such as the production of associated jets and the underlying event, are obtained from the original data event.

Tau embedding was introduced in CMS already during Run 1, when *double-tau embedding* was used to estimate the $Z/\gamma^* \to \tau_h\tau_h$ background using $\mu\mu$ events in searches for neutral Higgs bosons decaying to $\tau\tau$ [204], and *single-tau embedding* was applied for the estimation of top and electroweak backgrounds in the $H^{\pm} \to \tau^{\pm}\nu_{\tau}$ analyses [27, 195]. The ATLAS Collaboration has also used similar techniques in their corresponding analyses [26]. In Run 2, an improved version of the embedding technique has been introduced in the context of CMS $H \to \tau\tau$ analyses. Previously, the removal of muons and the addition of simulated tau leptons was performed at the level of particle flow candidates. In the new version the muons are cleaned from the events more carefully, by locating the hits and calorimeter energy deposits corresponding to the muon trajectory, and re-reconstructing the event without them. This new approach for double-tau embedding is described in Ref. [205] and applied in a $H \to \tau\tau$ search for the first time in Ref. [206]. In the following, the new approach is for the first time applied in the context of single-tau embedding for the $H^{\pm} \to \tau^{\pm}\nu_{\tau}$ channel.

The embedding procedure contains four separate steps. First, a suitable collection of single-muon events is selected (Section 12.2.1). Then the tracks and calorimeter energy deposits associated with the muon are cleaned from the event (Section 12.2.2).

Thirdly, a tau lepton decay is simulated, based on the kinematic properties of the muon (Section 12.2.3). Finally, the cleaned event and an event containing only the simulated tau lepton are merged into a hybrid event (Section 12.2.4). After a sample of embedded events is produced, it needs to be normalized as discussed in Section 12.2.5. In Section 12.2.6 we conclude by discussing the current status of the development of this method and its future prospects.

12.2.1 Event selection

Any event that contains a muon can be subject to the embedding procedure. However, since embedding is computationally expensive, it is desirable to select a sample of events that are likely to enter the signal region of the search in which the embedded sample is used for background estimation.

In case of the $H^\pm \to \tau^\pm \nu_\tau$ channel, we target a similar event topology as in the τ_h + jets final state, except that the events are required to contain a muon instead of a τ_h candidate. Thus the signal selection described in Chapter 11 is a natural starting point in designing the muon event selection for embedding. Obviously, it needs to be modified by replacing the signal trigger with a single-muon trigger and the τ_h selection with a similar muon selection. However, this approach has two limitations. Firstly, while $\vec{p}_T^{\text{miss}}$ is an important part of the signal event topology, it cannot be used in the selection of single-muon events, since the neutrino produced in the simulated τ_h decay change the $\vec{p}_T^{\text{miss}}$ of the event significantly during the embedding procedure. As angular selections also depend on $\vec{p}_T^{\text{miss}}$, they cannot be applied either until the embedding has been performed. Secondly, this approach implies that if the event selection is changed in any way, the embedded sample should be reproduced. While a looser selection of the single-muon events yields a larger sample to be processed, it allows changes in the signal selection (within the constraints set by this looser selection) and thus reduces the risk of multiple reprocessing efforts.

Taking these aspects into account, the following selection is designed to choose a sample of single-muon events to be embedded. A logical OR of three single-muon HLT paths is used: the first one requires a muon with $p_T > 24\,\text{GeV}$ and a loose track-based isolation, the second has the same p_T threshold but it uses both tracker and calorimeter information for isolation, and the third algorithm has a higher p_T threshold of $50\,\text{GeV}$ and no isolation requirement, which leads to improved trigger efficiency at high p_T.

As the p_T of the muon determines the p_T of the embedded tau lepton, and the signal selection requires the τ_h candidates to have $p_T > 50\,\text{GeV}$, only events containing a muon with offline $p_T > 45\,\text{GeV}$ are selected. Both trigger and offline selections restrict the muon pseudorapidity to $|\eta| < 2.4$. In order to select only prompt muons, tight muon identification criteria, as described in Section 8.2.1, are required. They include matching of the muon candidate to the primary vertex, efficiently suppressing secondary muons.

No isolation requirements are applied in the muon selection, except for the loose isolation requirements introduced by the trigger. Isolation requirements could bias the embedded sample, as the embedded events would always have a region with less hadronic activity, corresponding to the direction of the removed muon. To select only single-muon events, any additional muons passing the loose identification criteria and with $p_T > 20$, $|\eta| < 2.5$ are vetoed. To prevent unexpected correlations between the data and the background estimate obtained with the embedding method, the same events should not picked by both the signal selection and the selection of the single-muon events for embedding. This can be ensured by synchronizing the muon selection used in the embedding and the muon veto used in the signal selection such that any event containing a muon that is chosen for embedding is rejected in the signal selection.

At least 2 jets with $p_T > 25\,\text{GeV}$, $\eta < 5.0$ and $\Delta R(\mu, \text{jet}) > 0.1$ are required. No b jet identification requirement is applied, in order to produce an embedded sample that can be used in a flexible way with rapidly advancing b jet identification algorithms. One example of the selected events, picked from the data recorded in 2016, is shown in Figure 12.2 (left).

12.2.2 Cleaning

In the second step, the energy deposits from the selected muon, i.e. the hits in the tracker and muon systems and the clusters in calorimeters, are removed from the event.

The tracker and muon chamber hits associated with the fitted trajectory of the global muon are straightforward to identify and remove. In the case of calorimeter energy deposits, only the deposits attributed to the muon are removed, while keeping the hadronic activity close to the muon trajectory intact. The following approach is used. The muon trajectory is interpolated through the calorimeters. If the trajectory

Figure 12.2: Left: A selected single-muon event, containing one muon (red line) and three hadronic jets (green towers). Right: The same event after the muon is cleaned away and the event is reconstructed. The three hadronic jets (green towers) are not affected by the cleaning. Both pictures also show the ECAL and HCAL energy clusters (red and blue towers, respectively) and the $\vec{p}_T^{\text{miss}}$ (purple arrow).

intercepts a calorimeter cluster, the amount of energy corresponding to a minimum ionizing particle is subtracted from the cluster energy. If the energy falls below the noise threshold used in the PF clustering, the cluster is completely removed.

In principle, removal of the energy deposits can create nontrivial effects. For example, if a HCAL cluster is removed but an ECAL cluster is only reduced, the remaining ECAL energy is reconstructed as an additional photon in the cleaned event. In the opposite case, a neutral hadron would be reconstructed. This effect has been investigated, and its effect on the final results is found to be negligible, because the reconstructed extra particles have typically very low energies [205]. In Figure 12.2 (right), an exemplary single-muon event is shown as re-reconstructed after the cleaning step .

12.2.3 Tau simulation

The four-momentum of the cleaned muon is used to determine the four-momentum of the simulated tau lepton prior to its decay as follows. The three-momentum of the muon (p) is used as the three-momentum of the tau lepton. The energy of the tau lepton is set to $\sqrt{m_\tau^2 + p^2}$, where m_τ is set to $1.78\,\text{GeV}$.

The four-momentum of the tau is stored in a separate file in the LHE format [141], which describes a simple nonphysical event containing only the tau lepton. The polarization of the tau lepton is set to correspond to τ production in $W^\pm$ boson decay.

The LHE file is given as input to PYTHIA, which is used to simulate the tau decay, and then interfaced with GEANT4 for detector simulation. After the full simulation, the simulated event is reconstructed using the standard PF workflow. In all stages of simulation, only the single tau lepton is considered. The simulated events do not contain any additional activity, such as jets, underlying event or pileup, as these are already included in the cleaned event (to be merged with the simulated tau decay).

To minimize the statistical uncertainty in the background estimate from the embedded sample, it is desirable to maximize the number of embedded events that enter the signal region. While the tau embedding method can be used to produce embedded events with taus decaying via all possible decay channels, our goal is to estimate the background from events with a genuine tau lepton that decays hadronically. Thus the size of the embedded sample with τ_h candidates is increased by forcing the simulated tau lepton to always decay hadronically, and taking the corresponding branching fraction into account in the normalization.

To further increase the probability of the embedded events to enter the signal region, a feature of PYTHIA known as *kinematic filtering* is used to restrict the phase space of the simulated tau leptons so that they are likely to pass the signal selection. This also ensures that time and resources do not get wasted in computationally expensive detector simulation for events that have no chance of being used in the analysis. In the hadronic tau lepton decays, a large fraction of the initial energy of the tau lepton is carried away by a neutrino. Thus the visible hadronic decay products, which can be reconstructed as a τ_h candidate, have a lower p_T than the original tau lepton. Therefore despite the hard 45 GeV threshold for the muon p_T (equal to the p_T of the simulated tau), most τ_h candidates would not pass the stringent $p_T > 50$ GeV requirement used in the signal selection. Also the pseudorapidity range allowed for τ_h candidates ($|\eta| < 2.1$) is more narrow than the one used in muon selection ($|\eta| < 2.5$).

To account for the filtering in the normalization of the embedded sample, the tau decay is repeated 1000 times for each event. Only in some trials the decay products pass the kinematic filter. Only the last trial that passes the kinematic filter is stored, and the other ones are used to calculate the fraction of successful trials out of all 1000 trials. This fraction, which also accounts for only allowing hadronic decays, is stored as an event weight, which is taken into account in the normalization of the embedded sample as discussed later in Section 12.2.5.

Figure 12.3: Left: A selected single-muon event, containing one muon (red line) and three hadronic jets (green towers). Right: A hybrid event created by merging the cleaned event (shown in Figure 12.2, right) with a simulated τ_h decay. The τ_h is visible as a narrow hadronic jet (the narrow green tower pointing to southwest). Both pictures also shows the ECAL and HCAL energy clusters (red and blue towers, respectively). The $\vec{p}_\mathrm{T}^\mathrm{miss}$ (purple arrow) changes in the process, so it is recalculated for the hybrid event.

The usage of the three-momentum of the reconstructed muon to determine the three-momentum of the simulated tau lepton is based on two assumptions. Firstly, that the p_T spectra of taus (prior to decay) and muons are essentially identical, and secondly, that the p_T spectrum of the reconstructed muons is essentially identical to the true p_T distribution free of detector artifacts and reconstruction effects. The first assumption is reasonable, because the muons and tau leptons only differ in mass. As the muons and tau leptons considered here have $p_\mathrm{T} > 50\,\mathrm{GeV}$, the effect of the muon (tau lepton) mass of 0.1 (1.8) GeV on the kinematics is negligible. The second assumption was studied in Ref. [205]. The limited detector resolution was found to cause broadening of the muon p_T distribution, and the final state radiation from the selected muons to shift the reconstructed p_T distribution slightly downwards compared to the true distribution. The conclusion was however that these effects are negligible in case the embedded tau lepton decays hadronically.

12.2.4 Event merging

In the final step, the simulated tau decay event, obtained from the detector simulation, is combined with the original event from which the muon is cleaned. Ideally, this merging could be done at the level of hits and calorimeter deposits. In practice, the detector geometry is not perfectly described in simulations, so this type of merging

Figure 12.4: The effect of the embedding procedure on the jet multiplicity (left) and the p_T^{miss} distribution (right). The jet multiplicity remains essentially unchanged, as expected, while the p_T^{miss} distribution shifts (on average) towards higher values as the neutrinos from the simulated τ_h decays alter the value of $\vec{p}_T^{miss}$ in each embedded event.

would likely cause mismatches in positions of the τ_h energy deposits compared to the rest of the event, leading to residual effects and biases in the reconstructed embedded events. Thus the cleaned event and the simulated event are reconstructed separately, and then merging is performed at the level of reconstructed PF objects. This way a sample of hybrid events is created. In the end, the $\vec{p}_T^{miss}$ is recalculated for the merged event. The change in $\vec{p}_T^{miss}$ is illustrated in Figure 12.3, where the original single-muon event (left) is compared to the merged event (right).

In Figure 12.4, two distributions are compared between the original single-muon event (blue) and the embedded event with simulated τ_h decay products (red). On the left, the jet multiplicity distribution is shown (for reconstructed anti-k_T jets with distance parameter 0.4). This is an example of the distribution that is critical for the selection efficiency and modeled more accurately by the embedding method than by simulations. The figure confirms that the embedded events reproduce the jet distributions observed in the original single-muon events. On the right, the change in the p_T^{miss} distribution is illustrated. As the neutrino from a simulated τ_h decay contributes to the p_T^{miss}, it shifts towards higher values.

12.2.5 Normalization

After the embedded sample is produced, it is processed through the signal selection described in Chapter 11.

The embedded sample is used to estimate the background from the SM top or electroweak events with tau leptons originating from $W^{\pm}$ or Z decays. The total number of such *single-tau events* present in the data is denoted as N_{τ}. Prior to the offline selection, the normalization of the embedded sample should correspond to the number of single-tau events that pass the signal trigger:

$$N_{\tau}^{\text{trg}} = \epsilon_{\tau_{\text{h}}}^{\text{trg}} \, \epsilon_{p_{\text{T}}^{\text{miss}}}^{\text{trg}} \, N_{\tau}, \tag{12.1}$$

where $\epsilon_{\tau_{\text{h}}}^{\text{trg}}$ and $\epsilon_{p_{\text{T}}^{\text{miss}}}^{\text{trg}}$ are the selection efficiencies for the τ_{h} and $p_{\text{T}}^{\text{miss}}$ parts of the signal trigger described in Section 11.1.

However, the event selection described in Section 12.2.1 also picks up events that contain muons from tau lepton decays:

$$N_{\mu}^{\text{selected}} = N_{\text{prompt } \mu}^{\text{selected}} + N_{\tau \to \mu}^{\text{selected}}. \tag{12.2}$$

Thus if we define f to be the fraction of events with muons from tau decays, $f = N_{\tau \to \mu}^{\text{selected}} / N_{\mu}^{\text{selected}}$, we obtain

$$N_{\text{prompt } \mu}^{\text{selected}} = (1 - f) N_{\mu}^{\text{selected}}. \tag{12.3}$$

The single-muon trigger and selection efficiencies $\epsilon_{\mu}^{\text{trg}}$ and $\epsilon_{\mu}^{\text{sel}}$ relate the number of selected events to the total single-muon event count:

$$N_{\text{prompt } \mu}^{\text{selected}} = \epsilon_{\mu}^{\text{trg}} \epsilon_{\mu}^{\text{sel}} N_{\text{prompt } \mu}. \tag{12.4}$$

According to the lepton universality, the number of single-muon events with a muon produced in $W^{\pm}$ or Z boson decays is equal to the number of single-tau events, so

$$N_{\tau} = N_{\text{prompt } \mu} = \frac{N_{\text{prompt } \mu}^{\text{selected}}}{\epsilon_{\mu}^{\text{trg}} \, \epsilon_{\mu}^{\text{sel}}} = \frac{(1 - f) N_{\mu}^{\text{selected}}}{\epsilon_{\mu}^{\text{trg}} \, \epsilon_{\mu}^{\text{sel}}}. \tag{12.5}$$

Finally, we need to normalize the embedded sample according to Eq. (12.1), and include the event weights from the kinematic filter. If the (effective) integrated lu-

minosity of the final embedded sample (L_{embedded}), is different from the integrated luminosity of the data used in the analysis (L), this difference also needs to be included in the normalization. Difference in the luminosities can arise e.g. from different trigger prescales or because of computational constraints (only part of the muon events are embedded). Thus the final normalization of the embedded sample is

$$N_{\text{normalized}}^{\text{embedded}} = \frac{\epsilon_{\tau_{\text{h}}}^{\text{trg}} \, \epsilon_{p_{\text{T}}^{\text{miss}}}^{\text{trg}}}{\epsilon_{\mu}^{\text{trg}} \, \epsilon_{\mu}^{\text{sel}}} \frac{L}{L_{\text{embedded}}} (1 - f) \sum_i w_i^{\text{filter}} N_i^{\text{embedded}}. \tag{12.6}$$

12.2.6 Future prospects

The new version of the tau embedding method has been successfully implemented as described above, and preliminary studies using small samples of single-muon events have been performed to validate that each step (selection, cleaning, simulation and merging) produces the expected output.

The background estimate from embedding is subject to both statistical and systematic uncertainties. To minimize the statistical uncertainties, the full sample of single-muon events recorded during Run 2 (passing the criteria in Section 12.2.1) needs to be embedded. This computationally heavy effort is ongoing at the time of writing. Further optimization of the kinematic filtering used in the τ_{h} simulation (Section 12.2.3) can improve the selection efficiency for the embedded events, and hence further reduce the statistical uncertainty.

Additional validation of the method is still required, including extensive comparisons between embedded single-muon data, embedded simulated single-muon events, and the purely simulation-based background estimate. The amount of signal contamination needs to be quantified by injecting simulated signal events to the data samples used as input for embedding. While the signal contamination from $\mathcal{B}(\text{H}^{\pm} \to \mu^{\pm} \nu_{\mu,})$ events is expected to be negligible, a larger contamination can arise from events where the $\text{H}^{\pm}$ decays into $\tau^{\pm} \nu_{\tau}$ and the $\tau^{\pm}$ further decays into $\mu^{\pm} \nu_{\mu}$.

Finally, each term in Eq. (12.6) is subject to systematic uncertainties, which need to be propagated into the final background estimate. In the end, the comparison between the embedding estimate of the genuine-tau background, with full Run 2 single-muon data and with full systematic uncertainties, and the background estimate obtained from simulation, will determine which method provides the best performance.

12.3 Estimation of jet $\rightarrow \tau_h$ background from data

The background from events with jets misidentified as τ_h candidates (denoted as jet $\rightarrow \tau_h$) is estimated from a control region enriched in jets misidentified as τ_h and normalized by *transfer factors* (fake factors).

The background measurement is based on a so-called *ABCD method*, illustrated in Figure 12.5. To determine the jet $\rightarrow \tau_h$ background contribution in the signal region (D), a separate *control region* (C), free of signal contamination, is used to determine the shape of the background. To account for the different jet $\rightarrow \tau_h$ event yields in the two regions, these yields are compared between two other regions (A and B), deriving *transfer factors* that quantify the change in the background normalization from region A to region B. These transfer factors are then applied to normalize the background shape obtained from region C to the signal region D.

The measurement method is similar to the one used in previous CMS results on the $H^{\pm} \rightarrow \tau^{\pm} \nu_{\tau}$ channel, such as Ref. [27]. Among the jet $\rightarrow \tau_h$ events, QCD multijet production is the dominant process ($\approx$80% after all signal selection steps), compared to the electroweak/top events ($\approx$20%). The main difference between the old and new version of the background measurement is that previously only the QCD multijet background was estimated with the ABCD method, while the jet $\rightarrow \tau_h$ background contribution from electroweak/top was estimated directly from simulation. In the new version, the estimation obtained with the ABCD method accounts for both QCD multijet and electroweak/top components of the jet $\rightarrow \tau_h$ background.

This background measurement does not however consider isolated electrons or muons misidentified as τ_h, since the inverted identification requirement used to obtain the control sample rejects any isolated objects. The measurement of the background from isolated electrons or muons misidentified as τ_h is described in Section 12.4.

12.3.1 Selection of the control sample

The control region C enriched in jet $\rightarrow \tau_h$ events is obtained by inverting the MVA-based τ_h identification and isolation requirement used in the offline signal selection (Section 8.2.5). In other words, only events with very loosely isolated τ_h candidates (with the only isolation requirement coming from the trigger), failing the MVA-based selection, pass this *altered selection*.
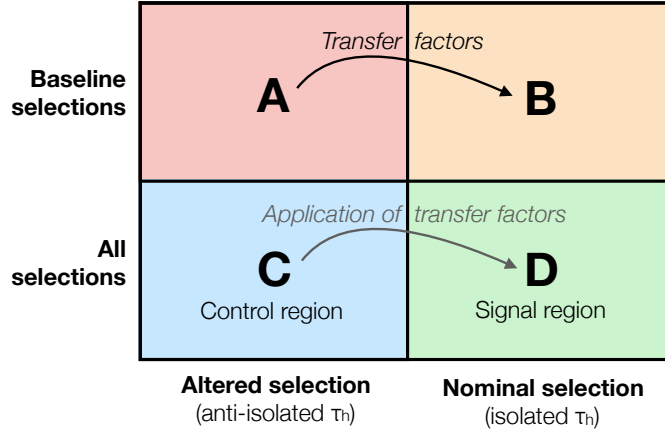
Figure 12.5: Illustration of the ABCD method: Transfer factors derived using events in regions A and B are used to normalize the background shape obtained from the control region C to the signal region D.

By construction, the control region C is dominated by jet $\to \tau_h$ events and it does not share any events with the signal region D.

The contamination of the control region from electroweak/top events containing only genuine τ_h candidates or isolated leptons misidentified as τ_h candidates ($e/\mu \to \tau_h$ events) is estimated from the simulation by processing the simulated top/electroweak background samples through the altered selection, normalizing these background processes to their theoretical cross sections and subtracting their contribution from the control sample:

$$N_{C,\,i}^{\mathrm{jet}\to\tau_h} = N_{C,i}^{\mathrm{data}} - N_{C,i}^{\tau\to\tau_h} - N_{C,i}^{e/\mu\to\tau_h}. \tag{12.7}$$

To account for the correlation between the p_T of the τ_h and p_T^{miss} as well as geometrical differences in detector response, the control sample is selected and normalized in separate bins of p_T (< 60, 60–80, 80–100 and $> 100\,\mathrm{GeV}$) and $|\eta|$ (< 0.6, 0.6–1.4 and > 1.4) of the τ_h candidate. In the above equations, i refers to one bin in p_T and $|\eta|$ of the τ_h candidate. In the following, we refer to these bins simply as the τ_h bins.

12.3.2 Normalization

The difference in the selection efficiency between the signal selection (*nominal selection*) and the altered selection is corrected by normalizing the control sample with suitable transfer factors.

As shown in Figure 12.5, the transfer factors are determined at an early point of event selections (regions A and B), after the *baseline selections* described in Section 11.2.2, and before applying the b jet identification, the offline selection on p_T^{miss} or the angular selections. After the baseline selections the number of selected events is still large and the risk for statistical fluctuations is small. In addition, a possible signal does not yet stand out from the large background yield, so there is no risk that it would get absorbed into the background estimate and bias the results.

The jets in the QCD multijet events and electroweak/top events are known to have different quark and gluon composition, implying different transfer factors [196]. Thus the transfer factors for misidentified τ_h from the QCD multijet events and from electroweak/top events are estimated separately and combined as a weighted average.

The transfer factor for the QCD multijet events is defined as

$$R_i^{\text{QCD}} \equiv \frac{N_{B,i}^{\text{QCD}}}{N_{A,i}^{\text{QCD}}}, \tag{12.8}$$

where $N_{B,i}^{\text{QCD}}$ ($N_{A,i}^{\text{QCD}}$) is the number of QCD multijet events passing the nominal (altered) baseline selections.

The $N_{A,i}^{\text{QCD}}$ yield is estimated by subtracting the simulated electroweak/top contribution (simulated events containing either genuine or non-genuine τ_h) from data. To estimate $N_{B,i}^{\text{QCD}}$, a binned maximum likelihood fit of p_T^{miss} templates to data is performed, using the fraction of the QCD multijet events as a fit parameter. The templates describe the expected shape of the p_T^{miss} distribution for each background component prior to the fit. The p_T^{miss} shape of the QCD multijet events is assumed to be similar in the phase space regions defined by the nominal and altered baseline selections, so the shape observed in region A is used as the fit template determining the QCD shape in region B.

From a set of common fit functions, the best description for the p_T^{miss} distribution of the total jet $\rightarrow \tau_h$ background, as well as its QCD multijet component, is achieved with a combined Rayleigh, Gaussian and exponential function

$$\frac{x - \mu_1}{\sigma_1^2} e^{\frac{-(x-\mu_1)^2}{2\sigma_1^2}} + \frac{e^{-(x-\mu_2)^2}}{\sigma_2\sqrt{2\pi}} + \mu_3 e^{-\sigma_3 x}, \tag{12.9}$$

where σ_i and μ_i are the fit parameters.

The template for electroweak/top events is obtained directly from simulation and fitted using the combined Gaussian and exponential function

$$\frac{e^{-(x-\mu_1)^2}}{\sigma_1\sqrt{2\pi}} + \mu_2 e^{-\sigma_2 x}. \tag{12.10}$$

All fits are performed separately in each p_T and $|\eta|$ bin. An example of the resulting fits is shown in Figure 12.6. The transfer factor for electroweak/top events is

$$R_i^{\text{EWK+top}} \equiv \frac{N_{\text{B},i}^{\text{EWK+top}}}{N_{\text{A},i}^{\text{EWK+top}}}. \tag{12.11}$$

It is estimated from simulation by counting the number of events $N_{\text{B},i}^{\text{EWK+top}}$ ($N_{\text{A},i}^{\text{EWK+top}}$) that do not contain a genuine tau lepton but still pass the baseline selections of the nominal (altered) selection.

Finally, the overall transfer factor for the control sample is determined in each τ_h bin as a weighted sum of the two process-specific transfer factors.

After the baseline selections, further cuts designed to suppress the QCD multijet events are applied. These cuts obviously alter the relative fraction of the QCD and electroweak/top components in the jet $\rightarrow \tau_h$ background. As the selection efficiencies and the transverse mass shape for the two fake tau background components are different, the final transfer factors must reflect the relative fraction of each component *after all selections*, even though the transfer factors R_i^{QCD} and $R_i^{\text{EWK+top}}$ are determined using the baseline selections.

Thus the final transfer factors R_i are defined as a weighted average, where the weight corresponds to the relative fractions of the QCD multijet and electroweak/top events in the control region after all selections:

$$R_i \equiv w_i R_i^{\text{QCD}} + (1 - w_i) R_i^{\text{EWK+top}}, \tag{12.12}$$

where w_i is the fraction of QCD events, with respect to all jet $\rightarrow \tau_h$ events in the control sample, estimated from the final event yields after all selections:

$$w_i = \frac{N_{\text{C},i}^{\text{QCD}}}{N_{\text{C},i}^{\text{jet} \rightarrow \tau_h}} = \frac{N_{\text{C},i}^{\text{data}} - N_{\text{C},i}^{\text{EWK+top}}}{N_{\text{C},i}^{\text{data}} - N_{\text{C},i}^{\tau \rightarrow \tau_h} - N_{\text{C},i}^{e/\mu \rightarrow \tau_h}}. \tag{12.13}$$

In summary, the measured number of jet $\to \tau_h$ background events expected in the signal region D is

$$\sum_i \left(N_{C,i}^{\text{data}} - N_{C,i}^{\tau \to \tau_h} - N_{C,i}^{e/\mu \to \tau_h}, \right) \times R_i \tag{12.14}$$

where $N_{C,i}^{\text{data}}$ is the number of events in data in the control region C and $N_{C,i}^{\tau \to \tau_h}$ ($N_{C,i}^{e/\mu \to \tau_h}$) is the number of simulated electroweak/top events with a τ_h candidate from a genuine tau lepton (from $e/\mu \to \tau_h$ misidentification), after all selection steps in the altered selection. The index i runs over all τ_h bins.

12.3.3 Systematic uncertainties

The jet $\to \tau_h$ background estimates obtained with the fake factor method described above are subject to three types of systematic uncertainties. Firstly, the simulated genuine-tau and $e/\mu \to \tau_h$ electroweak/top events subtracted from the control sample are subject to various systematic uncertainties. Secondly, the transfer factors have limited precision due to statistical uncertainties. Thirdly, while the method assumes that the m_T shape for the jet $\to \tau_h$ background is similar in the signal and control regions, statistical fluctuations can introduce differences between the two shapes.

The first group of uncertainties, i.e. the systematic uncertainties affecting the simulated samples used to estimate the genuine-tau and $e/\mu \to \tau_h$ contribution which is subtracted from data, are propagated through the background estimation and scaled down by the fraction of these simulated events in the control region. They include e.g. uncertainties related to the trigger, identification of physics objects, and theoretical cross sections, as discussed in detail in Chapter 13.

Because these uncertainties alter events that are *subtracted* from data in the jet $\to \tau_h$ background estimate, each of them is *anti-correlated* with respect its counterparts that affect other processes. For example, when the uncertainty in the integrated luminosity (Section 13.8) that affects all simulated processes is varied *upwards* by 2.5% for those processes, it is simultaneously varied *downwards* for the jet $\to \tau_h$ background, and since it only affects the normalization of subtracted simulated events, the outcome is a variation of -0.4%.
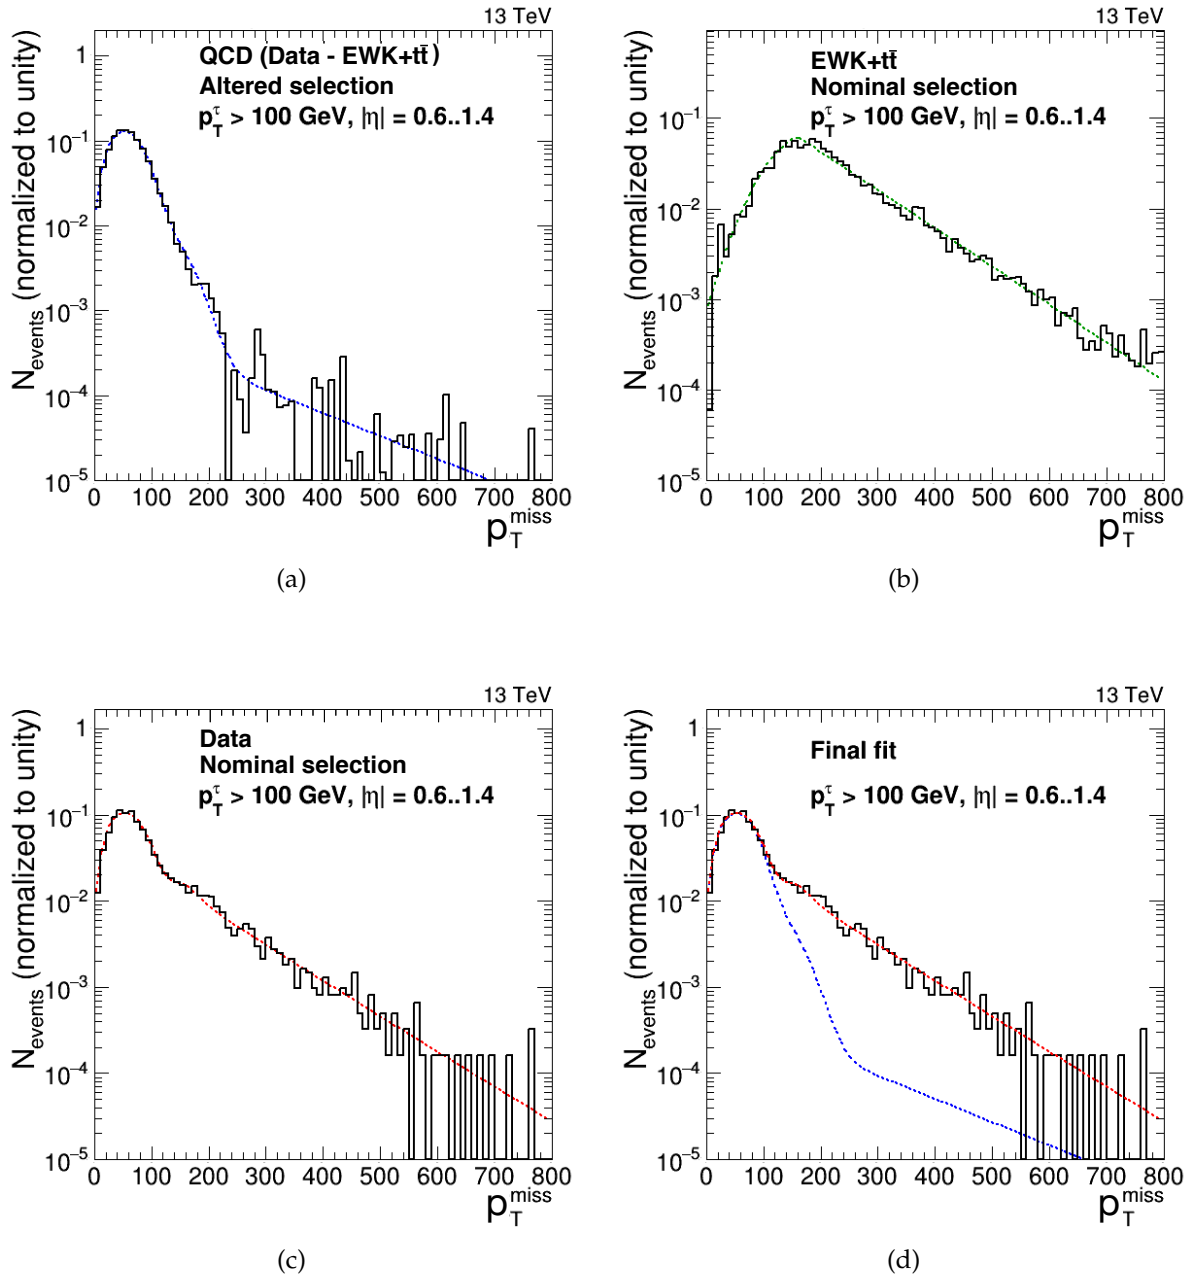
Figure 12.6: Distribution of p_T^{miss} and the corresponding fitted shapes for: **(a)** QCD multijet events from the altered selection (electroweak/top subtracted from data), **(b)** electroweak/top from the nominal selection and **(c)** data from the nominal selection. All these templates are used to perform the final fit **(d)**, where the QCD shape taken from **a** (blue line) is fitted together with electroweak/top shape **b** to data **c** (red line), yielding an estimate for the fraction of QCD events in the signal region. The fits are performed separately in different bins of p_T and $|\eta|$ of the τ_h candidate. Here the bin with $p_T > 100$ GeV and $0.6 < |\eta| < 1.4$ is shown.

The uncertainties in transfer factors for each τ_h bin (R_i) are calculated by propagating the statistical uncertainties of R_i^QCD, $R_i^\text{EWK+top}$ and w_i through Eq. 12.12. The effect on the overall normalization is then obtained by combining all the R_i uncertainties, treating them as uncorrelated (as they are of statistical origin).

The uncertainty from shape differences of the m_T distributions between the signal and control regions is determined as follows. First, the m_T distributions for the jet $\rightarrow \tau_\text{h}$ events are obtained for both the signal and control samples (after all selections) and normalized to unity. Then for each bin in the distributions, the statistical uncertainty due to the limited number of events in the bin is calculated. The normalized signal region distribution is divided by the normalized control region distribution, and the statistical uncertainty of the resulting quotient distribution is calculated by error propagation and applied as a shape uncertainty.

12.3.4 Validation

The jet $\rightarrow \tau_\text{h}$ background measurement method is validated in several ways. Firstly, the chosen binning, based on the τ_h p_T and $|\eta|$, might bias the results. To ensure that this is not the case, the measurement is performed with several alternative binning schemes, including a measurement inclusive in $|\eta|$ and a measurement inclusive both in the τ_h p_T and $|\eta|$. The resulting background estimates are found to be consistent within the uncertainties.

Secondly, a test is performed to ensure that no significant bias arises from deriving the transfer factors at an early stage of an event selection, and applying them to normalize the jet $\rightarrow \tau_\text{h}$ background after all selections. While the change in the relative fractions of the QCD multijet and electroweak/top components is accounted for by Eq. (12.12), the quark/gluon composition of the QCD multijet component itself can also change as additional selection steps such as b-tagging are applied. This would alter R_i^QCD in a way that is not accounted for by our method. To quantify this effect, the transfer factors derived after the baseline selections are compared to another set of transfer factors, obtained after applying both the baseline selections and the b jet selection criteria. The two sets of transfer factors are found to agree within the statistical uncertainties, indicating that the chosen approach for determining the transfer factors is reasonable.

Thirdly, a key assumption of this fake factor method is that the shape of the m_T distribution is very similar in the signal region D and control region C, so that it is sufficient to correct the overall normalization with the transfer factors (as a function of

Figure 12.7: Comparison of jet $\to \tau_h$ transverse mass distributions between events with τ_h candidates that pass the isolation requirements ("Isolated"), and events with τ_h candidates that fail them ("Anti-isolated"), in the $R_\tau > 0.75$ category (left) and $R_\tau < 0.75$ category (right). To suppress possible signal contamination, a veto on b-tagged jets is applied, contrary to the signal selection. The ratio of the two normalized shapes is also shown, together with statistical uncertainties propagated to the ratio. In both categories, the shapes agree within their statistical uncertainties.

τ_h p_T and $|\eta|$). To verify this assumption, the method is used to estimate the jet $\to \tau_h$ background in a *validation region* that is similar to the signal region but depleted from $H^\pm$ signal. The validation region is defined identically to the signal region, except that the b jet identification criterion is inverted, i.e. a b jet veto is applied to suppress the signal.

In the validation region, the shapes of the jet $\to \tau_h$ background are compared between the events that pass the τ_h isolation requirements (similar to the nominal selection), and events with τ_h candidates that fail them (similar to the altered selection). The shapes are obtained subtracting the contributions from simulated genuine-tau and $e/\mu \to \tau_h$ contributions from the data. The comparison is shown in Figure 12.7, separately for each R_τ category. In both categories, the shapes of the distributions are found to agree within their statistical uncertainties.

12.4 Estimation of e/$\mu \to \tau_h$ background from simulation

The backgrounds from top and electroweak events containing an electron or a muon misidentified as τ_h are estimated from simulation using the same samples and cross sections as in the case of the genuine-tau background (Section 12.1).

The only difference is that instead of matching the reconstructed τ_h candidates to generator-level τ_h decay products, they are matched to a generator-level electron or muon using a ΔR cone of 0.1. As the τ_h identification efficiency is different for genuine taus and leptons misidentified as τ_h, the corrections to the efficiency and the related uncertainties are also different. They are discussed later in Section 13.1.2.

The e/$\mu \to \tau_h$ background arises mostly from $t\bar{t}$ events. After the τ_h selection and before any other offline selections, the electrons misidentified as τ_h constitute 2.3% of the selected simulated $t\bar{t}$ events, the muons misidentified as τ_h contribute 1.8%, while the remaining 95.9% are events with genuine taus. After all selections, the total contribution from the e/$\mu \to \tau_h$ background is less than 2% of the total background.

12.5 Background selection efficiency

The final background yield estimates are summarized in Figure 12.8, showing the overall background event yields as a function of subsequent selection requirements, compared with the event yields observed in data, in the $R_\tau > 0.75$ (left) and $R_\tau < 0.75$ (right) categories. The estimates are found to agree with the observed data across the selection steps.

For the background processes estimated from simulation, namely the top and electroweak events with genuine tau leptons (Section 12.1) and electrons or muons misidentified as τ_h (Section 12.4), the effect of each selection step is detailed in Table 12.1. The yields are tabulated for the $R_\tau > 0.75$ category. In the event yields of Figure 12.8 and Table 12.1, all corrections applied to simulated samples are taken into account.

Figure 12.8: The estimated number of background events from different background processes as a function of the different selection steps, compared to the yields observed in data. The background estimates are found to agree with the observed data across the selection steps.

Table 12.1: Estimated number of events from genuine-tau and $e/\mu \to \tau_h$ backgrounds, as a function of the applied selections and including scale factors (SF). The event yields for each production process are normalized according to the theoretical cross section for that process. Tau selection includes the $R_\tau > 0.75$ selection. The statistical uncertainties are also shown.

Process	$t\bar{t}$	W+jets	Single top	Z/γ^*	Diboson
All events	33366800 ± 3500	135126000 ± 20000	10776800 ± 1000	16197500 ± 3200	6377500 ± 1700
Trigger	1019460 ± 570	1507716 ± 800	139254 ± 140	346262 ± 360	59503 ± 130
Tau selection	7000 ± 43	12222 ± 58	1225 ± 14	3219 ± 27	708 ± 14
Trigger & tau id. SF	5374 ± 32	9631 ± 44	969 ± 11	2338 ± 19	555 ± 11
Lepton veto	3747 ± 26	9190 ± 42	744 ± 10	1479 ± 15	397.7 ± 9.3
≥ 3 hadronic jets	3353 ± 24	5519 ± 30	602.5 ± 8.7	935 ± 10	269.7 ± 7.7
≥ 1 b-tagged jets	2487 ± 21	624.3 ± 9.4	419.5 ± 7.2	127.1 ± 3.4	37.9 ± 2.9
b-tagging SF	2430.1 ± 20.1	632.6 ± 9.4	411.7 ± 7.1	129.7 ± 3.4	38.3 ± 3.0
$p_T^{miss} > 90 \, GeV$	2289.8 ± 19.5	605.9 ± 9.3	392.7 ± 7.0	106.7 ± 3.1	36.2 ± 2.9
Angular selections	2253.0 ± 19.3	604.6 ± 9.3	390.4 ± 7.0	104.7 ± 3.0	35.6 ± 2.9
Final event yield	2253.0 ± 19.3	604.6 ± 9.3	390.4 ± 7.0	104.7 ± 3.0	35.6 ± 2.9

Chapter 13

Corrections and systematic uncertainties

Even though the simulation of production and decay processes as well as the detector response is carried out in great detail, it is not perfect. Thus the data and simulation are carefully compared to find imperfections in the modeling that might affect the event reconstruction and selection techniques and thereby bias the analysis results. When such discrepancies are found, they are corrected, and the limited precision of the methods is taken into account by assigning systematic uncertainties to the signal and background estimates obtained from simulation. Also the data-driven estimation of the jet $\to \tau_h$ background is subject to systematic uncertainties. This chapter describes the corrections and systematic uncertainties included in the analysis. A summary of these uncertainties is given in Table 13.1.

All uncertainties are incorporated in the statistical analysis as nuisance parameters as described in Chapter 9. Some of them are normalization uncertainties affecting only the final event yield for a process, whereas others also modify the shape of the final m_T distributions. The uncertainties from different sources are assumed to be uncorrelated. Each source of uncertainty is treated as fully correlated among the signal and background processes, except for the few special cases mentioned in the following.

13.1 Selection efficiencies

13.1.1 Trigger efficiency

The trigger efficiencies for the τ_h and $p_T^{\text{miss,calo}}$ parts of the trigger are measured separately, as described in Section 11.1.1. The results are used to derive corrections for the simulated events and estimate the related systematic uncertainties.

Table 13.1: Effect of systematic uncertainties on the final event yields in %, prior to the fit, summed over both R_τ categories. For the H$^\pm$ signal, the values corresponding to $m_{\mathrm{H}^\pm} = 200\,\mathrm{GeV}$ are shown. The "shape" uncertainties can modify both the shape and the normalization of the final m_T distributions, as opposed to the other uncertainties that only affect the normalization.

Source	Shape	H$^\pm$ (200 GeV)	Jets $\to \tau_h$	$t\bar{t}$	Single t	Electroweak
$\tau_h + p_T^{\mathrm{miss}}$ trigger efficiency	✓	10.3	2.0	9.0	7.3	7.4
τ_h identification	✓	4.2	0.6	4.0	4.0	4.1
Lepton veto efficiency		0.3	—	0.4	0.3	< 0.1
Jet energy scale and resolution	✓	4.8	0.4	2.0	2.2	4.3
τ_h energy scale	✓	1.8	0.6	2.2	2.1	2.3
Unclustered p_T^{miss} energy scale	✓	0.6	< 0.1	< 0.1	< 0.1	< 0.1
b jet identification	✓	3.5	0.8	2.9	2.9	5.9
Integrated luminosity		2.5	0.4	2.5	2.5	2.5
Pileup	✓	0.6	< 0.1	< 0.1	< 0.1	< 0.1
Jets misid. as τ_h estimation	✓	—	6.1	—	—	—
Cross section (scales, PDF)		—	0.8	5.5	5.3	3.3
Top quark mass		—	0.4	2.8	2.2	—
Acceptance (scales, PDF)		5.1	0.5	2.8	2.8	6.8
Total		14.0	6.6	12.9	11.6	13.9

To correct the simulated events to match the efficiencies measured in the data, scale factors are defined as ratios of fitted efficiencies for data (ϵ_{data}) and simulation (ϵ_{sim}). The corrections are applied to the simulated samples as a function of the τ_h p_T and p_T^{miss}.

The fit uncertainties are used to estimate the uncertainty in the scale factor $\epsilon_{\mathrm{data}}/\epsilon_{\mathrm{sim}}$ as follows. Both ϵ_{data} and ϵ_{sim} are varied up and down independently, resulting in four variations for both the τ_h and p_T^{miss} parts of the trigger: $\epsilon_{\mathrm{data}}^{\mathrm{up}}/\epsilon_{\mathrm{sim}}$, $\epsilon_{\mathrm{data}}^{\mathrm{down}}/\epsilon_{\mathrm{sim}}$, $\epsilon_{\mathrm{data}}/\epsilon_{\mathrm{sim}}^{\mathrm{up}}$ $\epsilon_{\mathrm{data}}/\epsilon_{\mathrm{sim}}^{\mathrm{down}}$. The effect of each variation is propagated into the final m_T distributions, and the varied distributions are applied as four independent nuisance parameters, corresponding to variations of ϵ_{data} and ϵ_{sim} for both the τ_h and p_T^{miss} parts of the trigger, such that each nuisance contains both up and down varied m_T templates.

13.1.2 τ_h identification and isolation

The agreement of the τ_h identification efficiency between data and simulated samples is studied using the tag-and-probe technique [175].

While the efficiencies agree within statistical uncertainties, a small difference in the central value is observed and the simulated events are corrected for it by applying a scale factor of 0.99 for all simulated events with τ_h candidates from genuine tau leptons. The uncertainty in the measurement is 5% and it is applied as a normalization uncertainty.

For the τ_h candidates with large transverse momentum, an additional uncertainty of $^{+5}_{-35}\%\,p_T/\,\text{TeV}$ is applied as a shape uncertainty to account for the uncertainty in the extrapolation of the measured efficiencies to the high-p_T region, where not enough data is available for a precise estimate. This extrapolation is performed with polynomial fits as detailed in Ref. [175].

The small background contribution from events with isolated electrons or muons misidentified as τ_h is estimated from simulation as described in Section 12.4. To reproduce the misidentification rates measured in data, these events are corrected with scale factors derived in Ref. [175], ranging from 1.40 ± 0.12 to 1.90 ± 0.30 (1.12 ± 0.04 to 2.39 ± 0.16) for electrons (muons), such that the scale factors and their uncertainties increase with $|\eta|$. The uncertainties are propagated to m_T distributions and incorporated as shape uncertainties, independently for electrons and for muons misidentified as τ_h.

13.1.3 Lepton identification and isolation

For the veto of isolated electrons or muons, the uncertainty from the identification and isolation is calculated as

$$\frac{N^{\text{vetoed}}}{N^{\text{selected}}} \times \Delta_{\text{ID}}, \tag{13.1}$$

where N^{vetoed} (N^{selected}) corresponds to the number of events failing (passing) the lepton veto step, where events with identified and isolated electrons/muons are rejected. For electrons (muons), the uncertainty in identification and isolation efficiency Δ_{ID} is 1% (2%). The veto uncertainties for electrons and muons are applied as two independent normalization uncertainties.

13.1.4 B jet identification

The differences in b-tagging and mistagging efficiencies between data and simulation are corrected by applying scale factors to simulated events [171]. The correction affects the final m_T shapes, so the related uncertainties are considered as shape uncertainties.

The *per-event scale factor* used in the correction is calculated for each simulated event separately taking into account the jet composition of the event and the corresponding *per-jet scale factors*, determined as a function of the p_T, η and jet flavor.

We refer to jet flavors as b, c, g, and u/d/s, where g refers to gluons while other letters refer to the quark origin of the jet (as determined from simulation). In this analysis, the c jets are treated as b jets, and the dependence of the correction on jet η is found to be small, so it is neglected.

The difference in b-(mis)tagging efficiency of a jet is quantified by the per–jet scale factors, f_{tag} and f_{mistag}, defined as

$$f_{tag}(p_T) = \frac{\varepsilon_{tag}^{data}(p_T)}{\varepsilon_{tag}^{simulation}(p_T)}, \quad f_{mistag}(p_T) = \frac{\varepsilon_{mistag}^{data}(p_T)}{\varepsilon_{mistag}^{simulation}(p_T)}, \tag{13.2}$$

where ε_{tag} (ε_{mistag}) is the b-(mis)tagging efficiency. The b-(mis)tagging efficiencies of jets are determined using simulated $t\bar{t}$ events, with all baseline selection steps applied as described in Section 11.2.2.

The probability P for the event to pass the b jet identification step is

$$P = \prod_{i=1}^{N_{b,c\ tagged}} \varepsilon_{tag,i} \prod_{j=1}^{N_{b,c\ not\ tagged}} (1 - \varepsilon_{tag,j})$$
$$\times \prod_{k=1}^{N_{uds,g\ tagged}} \varepsilon_{mistag,k} \prod_{l=1}^{N_{uds,g\ not\ tagged}} (1 - \varepsilon_{mistag,l}). \tag{13.3}$$

The per–event scale factor quantifies the difference in this probability between data and simulation:

$$SF = \frac{P(\text{data})}{P(\text{simulation})}. \tag{13.4}$$

Using Equations (13.2) and (13.3), the per–event scale factor (SF) used in the correction can be written in terms of per–jet scale factors ε and b-(mis)tagging efficiencies f as:

$$
\begin{aligned}
\text{SF} = {} & \prod_{i}^{N_{\text{b,c tagged}}} f_{\text{tag},i} \prod_{j}^{N_{\text{b,c not tagged}}} \left(\frac{\frac{1-f_{\text{tag},j}}{\varepsilon_{\text{b},j}}}{1 - \varepsilon_{\text{b,c},j}} \right) \\
& \times \prod_{k}^{N_{\text{uds,g tagged}}} f_{\text{mistag},k} \prod_{l}^{N_{\text{uds,g not tagged}}} \left(\frac{\frac{1-f_{\text{mistag},l}}{\varepsilon_{\text{uds,g},l}}}{1 - \varepsilon_{\text{uds,g},l}} \right) .
\end{aligned}
\tag{13.5}
$$

The uncertainties in the per–event b-tagging and mistagging scale factors incorporate the uncertainties in the per–jet scale factors as well as the measured b-(mis)tagging efficiencies, propagated through Eq. (13.5). In the error propagation, the per–jet scale factors f_{tag} and f_{mistag} and the measured efficiencies are assumed to be uncorrelated. The uncertainties for the c jets are conservatively taken to be twice as large as the b jet uncertainties. Finally, the b-tagging and mistagging uncertainties are treated as uncorrelated and thus included in the analysis as two separate shape nuisances.

13.2 Energy scales

The systematic uncertainties related to the energy measurements and calibrations for jets, τ_h candidates and p_T^{miss} are taken into account as shape uncertainties by performing variations in the energy scales, reprocessing the events through the event selection and using the resulting varied m_T distributions to define the shape nuisances.

13.2.1 Jet energy scale and resolution

The three-step jet energy scale calibration is described in Section 8.2.3. The systematic uncertainty related to this calibration procedure is estimated by propagating all (dozens of) uncertainties that affect the calibration through the full calibration chain, taking into account the correlations between the uncertainties, as detailed in Ref. [167]. As a result, the uncertainty is parameterized as a set of asymmetric factors that vary the jet momenta up and down as a function of the jet p_T and η. These variations are then propagated to the final m_T distributions and taken into account as a shape uncertainty.

Comparisons of data and simulated events show that the jet energy resolution in data is inferior to the resolution in simulated events. Therefore additional smearing is applied

to the jets in simulated samples, by scaling their four-momenta by p_{T}-dependent correction factors. The systematic uncertainty from jet energy resolution is derived in a dedicated study [167], where different uncertainties that affect the correction factors are considered and total uncertainties for the correction factors are derived. Similarly to the jet energy corrections, a shape uncertainty is obtained from these results by varying the jet resolution corrections by their uncertainties and propagating the changes to the final m_{T} distributions.

13.2.2 τ_{h} energy scale

A correction to the energy scale of the τ_{h} candidates is derived by comparing the $e\tau_{\mathrm{h}}$ and $\mu\tau_{\mathrm{h}}$ final states of $Z/\gamma^* \to \tau^+\tau^-$ events between data and simulation [175]. The τ_{h} energy scale is determined by maximum-likelihood fits based on the visible mass of the τ_{h} candidate and the mass of the $\ell\tau_{\mathrm{h}}$ system, separately for each τ_{h} decay mode.

This study results in corrections ranging from 0.995 to 1.011 depending on the τ_{h} decay mode. The corrections are applied in simulated events containing τ_{h} candidates with energies up to 400 GeV, with an uncertainty of $\pm 1.2\%$. Above 400 GeV, no correction is applied, but a larger $\pm 3\%$ uncertainty is applied.

13.2.3 $p_{\mathrm{T}}^{\mathrm{miss}}$ unclustered energy scale

The variations of the jet energy scale and resolution are propagated to the type-I corrected $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$, calculated as described in Section 8.2.6.

The uncertainties arising from the unclustered energy deposited in the detector are also included. The uncertainty in the unclustered energy is evaluated based on the measured momentum resolutions for the PF candidates in the event. Typically the largest contributions arise from neutral hadrons, reconstructed mostly in the HCAL, and from the particles reconstructed in the HF [165].

13.3 Jet $\to \tau_{\mathrm{h}}$ background estimation

The uncertainties related to the jet $\to \tau_{\mathrm{h}}$ background measurement are included in the analysis as described in Section 12.3.3.

13.4 Cross sections

The cross sections used for the initial normalization of each simulated background process are subject to theoretical uncertainties related to the choice of renormalization and factorization (RF) scales and PDFs [207]. For $t\bar{t}$ and single top quark processes, also the uncertainty in m_t must be taken into account. It is estimated by varying m_t by 1.0 GeV around the nominal value of 172.5 GeV.

The cross section uncertainties are taken into account as separate normalization uncertainties for each background process. For the dominating $t\bar{t}$ background, the total cross section uncertainty is approximately $^{+6.2}_{-6.6}$% at $\sqrt{s} = 13$ TeV. The total uncertainties are smaller for other simulated background processes such as W+jets (3.8%) or Z/γ^* ($^{+4.0}_{-3.8}$%).

13.5 Acceptance uncertainties

The theoretical uncertainties related to parton distribution functions (PDFs) and the choice of renormalization and factorization (RF) scales affect not only the cross sections for the simulated processes, but also the selection efficiency for the simulated events and hence the final event yields. We refer to these selection efficiency uncertainties as *acceptance uncertainties*.

For acceptance uncertainties related to the RF scales, these two scales are varied by factors of 0.5 and 2, excluding the extreme variations where one scale is varied by 0.5 and the other one by 2. The envelope of six variations is used to determine the total uncertainty.

As the varied m_T distributions are found to have a similar shape as the nominal one, the RF acceptance uncertainties are applied as normalization uncertainties. The uncertainty is 4.8% for signal samples up to 750 GeV and 1.2% for larger $H^\pm$ masses, 2% for $t\bar{t}$ and single top backgrounds and 5% for other simulated backgrounds.

For the PDF acceptance uncertainty, 100 simulated replicas for each sample are drawn from the PDF probability distribution, and the standard deviation of these weights is taken as the PDF uncertainty. Also these uncertainties are applied as normalization uncertainties: -0.4/+1.7% for signal samples, -2.0/+0.27% for $t\bar{t}$ and single top backgrounds and -3.3/+4.6% for other simulated backgrounds.

When a common PDF or RF acceptance uncertainty value is used for several processes (e.g. for $t\bar{t}$ and single top processes), the uncertainties for those processes are treated as fully correlated.

13.6 Pileup modeling

As discussed in Section 4.3.1, the average number of simultaneous proton–proton interactions (pileup) depends on the instantaneous luminosity in each bunch crossing as delivered by the LHC. The amount of pileup in data is estimated by multiplying the measured instantaneous luminosity by the total inelastic cross section of 69.2 mb. In simulation, the pileup for each event is sampled from a predefined distribution, so the simulated events need to be reweighed to make their pileup distribution match with the one observed in data.

The systematic uncertainty in the pileup modeling is estimated by shifting the total inelastic pp production cross section by $\pm 5\%$ around its nominal value [208], and repeating the reweighing procedure. The effect of variations is propagated to the final m_T distributions as a shape uncertainty.

13.7 Signal modeling

All uncertainties listed in Sections 13.1, 13.2, 13.5 and 13.6 are applied not only for the relevant simulated backgrounds, but also for all $H^{\pm}$ signal samples. For $H^{\pm}$ signal samples up to $m_{H^{\pm}} = 165$ GeV, all uncertainties related $t\bar{t}$ background are also applied (and fully correlated with uncertainties of the $t\bar{t}$ background), as the light $H^{\pm}$ is assumed to be dominantly produced in a decay of a top quark.

When the simulated signal samples for different $H^{\pm}$ mass points are compared, it is found that events from the intermediate-mass samples have higher selection efficiency and thus larger final event yields compared to the other signal samples. In principle, three effects contribute to this difference: Firstly, different production processes dominate for different signal mass hypotheses as discussed in Section 3.2, yielding different kinematic distributions that change the selection efficiency. Secondly, in the intermediate region the different production processes can interfere significantly. Thirdly, even though the same event generator software is used for all signal samples, the intermediate signal samples are simulated at LO while the light and heavy samples are based on an NLO model.
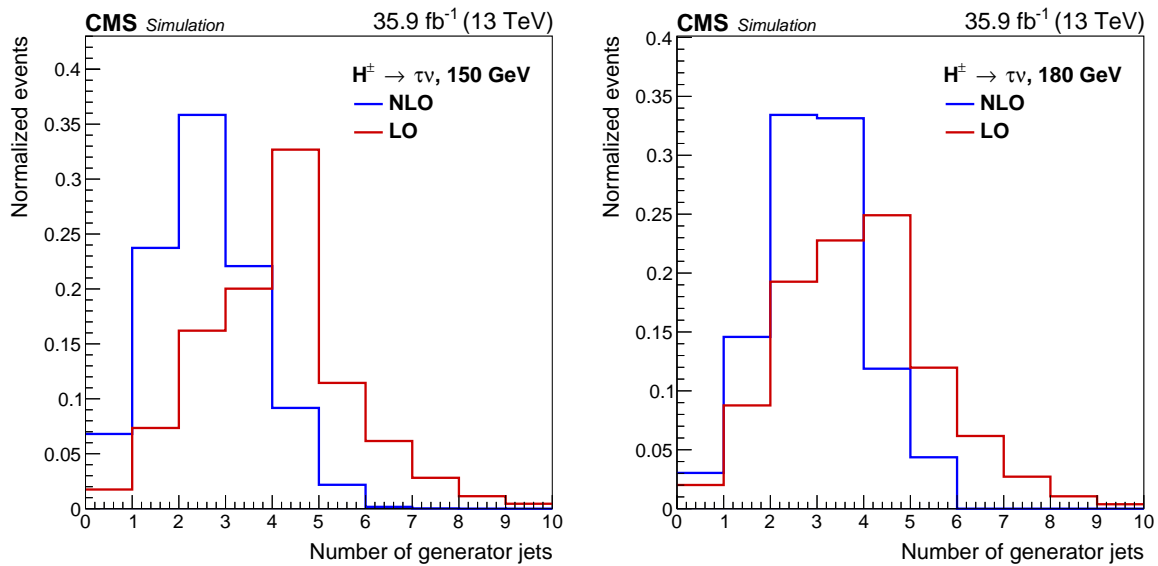
Figure 13.1: Generator level jet multiplicity distributions for next-to-leading order (NLO, blue curve) and leading order (LO, red curve) simulated samples for a charged Higgs boson with $m_{H^\pm} = 150\,\text{GeV}$ (left) and $m_{H^\pm} = 180\,\text{GeV}$ (right), decaying to $\tau^\pm \nu_\tau$ (inclusive in τ decay modes). The $H^\pm$ production in association with a top quark is assumed and the four-flavor scheme is used in the simulation.

While the two first effects are genuine physical behavior of the events, the third one results from shortcomings of the chosen signal modeling techniques. In order to disentangle these effects, a set of LO signal samples with light $H^\pm$ production (140, 150, 155, 160 GeV) and with heavy $H^\pm$ production (180, 200, 220 GeV) were produced. By definition these samples do not model the intermediate mass range effects, so the third effect is the only one to contribute.

A difference in the selection efficiency similar to the intermediate region is observed, so it can be concluded that the third effect is dominant and needs to be corrected. The difference in selection efficiency is found to originate mostly from different jet multiplicity distributions between the LO and NLO samples, shown in Figure 13.1.

The LO-to-NLO correction for the LO intermediate-range samples is derived by comparing the final m_T distributions between the LO and NLO signal samples below (140–160 GeV) and above (180–220 GeV) the intermediate mass range. As the shapes of the m_T distributions are found to agree, the LO-to-NLO correction is performed with scale factors that alter the final event yield for the LO samples.

The NLO/LO signal yield ratio is found to be approximately constant in the light region ($H^\pm < m_t$) and in the heavy region ($H^\pm > m_t$). The ratios in $R_\tau < 0.75$ and

$R_\tau > 0.75$ categories are found to agree within statistical uncertainties. Thus the scale factors are derived by calculating the ratios with inclusive R_τ selection and taking the average of the ratios in the light and heavy $H^\pm$ mass regions separately. These scale factors are applied to all intermediate-mass samples used in the analysis.

The scale factor derived from 140–160 GeV light $H^\pm$ samples (0.41 ± 0.12) is applied for intermediate mass points below $m_t - m_b$ mass (from 145 to 165 GeV), while the scale factor derived from 180–220 GeV samples (0.65 ± 0.13) is applied for intermediate-mass samples above $m_t - m_b$ (from 170 to 200 GeV).

The overall effect of the correction is to scale down the signal event yield, resulting in more conservative results than obtained by using the LO samples without this correction. As discussed in Section 10.2, for the final results the LO intermediate-range samples are used only in the 165–175 GeV range, and NLO samples are used elsewhere.

The systematic uncertainty related to this correction is calculated from the statistical uncertainties of the ratios, propagating them through the averaging used to calculate the scale factor. It is taken into account as a normalization uncertainty for the 165–175 GeV intermediate-range mass points included in the final results.

13.8 Luminosity measurement

The uncertainty in the measurement of the integrated luminosity is 2.5%, as estimated in Ref. [209]. This normalization uncertainty affects equally the normalization of all simulated background processes.

Chapter 14

Results

In this chapter, the final results of the data analysis described in Chapters 10–13 are presented. First, the transverse mass distributions observed in data are compared to those obtained from background estimations and the $H^{\pm}$ signal hypotheses. Then, model-independent exclusion limits are calculated using the m_T distributions. Finally, the results are combined with those from the analysis targeting the leptonic final states of the $H^{\pm} \to \tau^{\pm} \nu_\tau$ decay channel. The combined limits are presented and interpreted in the context of the MSSM.

14.1 Transverse mass distributions

After all selections, the transverse mass defined by Eq. (10.1), is calculated using the $\vec{p}_T$ of the reconstructed τ_h candidate and $\vec{p}_T^{\text{miss}}$. The transverse mass distributions are formed separately in the two R_τ categories. They are binned according to the statistical precision of the samples describing different background processes, ensuring that all bins are populated by background events. This leads to wider bins in the tail of the distribution. To prevent any overflow, the last bin is extended up to 5 TeV.

To study the agreement between the background estimates and the observed data, a binned maximum likelihood fit is performed simultaneously in the two categories under the background-only hypothesis. The systematic uncertainties are incorporated as nuisance parameters in the likelihood and profiled in the fit according to their probability density functions, taking correlations into account. For normalization uncertainties, log-normal probability density functions are used. For shape uncertainties, horizontal morphing [182] is used to derive continuous probability densities from the nominal and varied m_T shape templates. The statistical uncertainties are included using the Barlow-Beeston-lite approach [183].
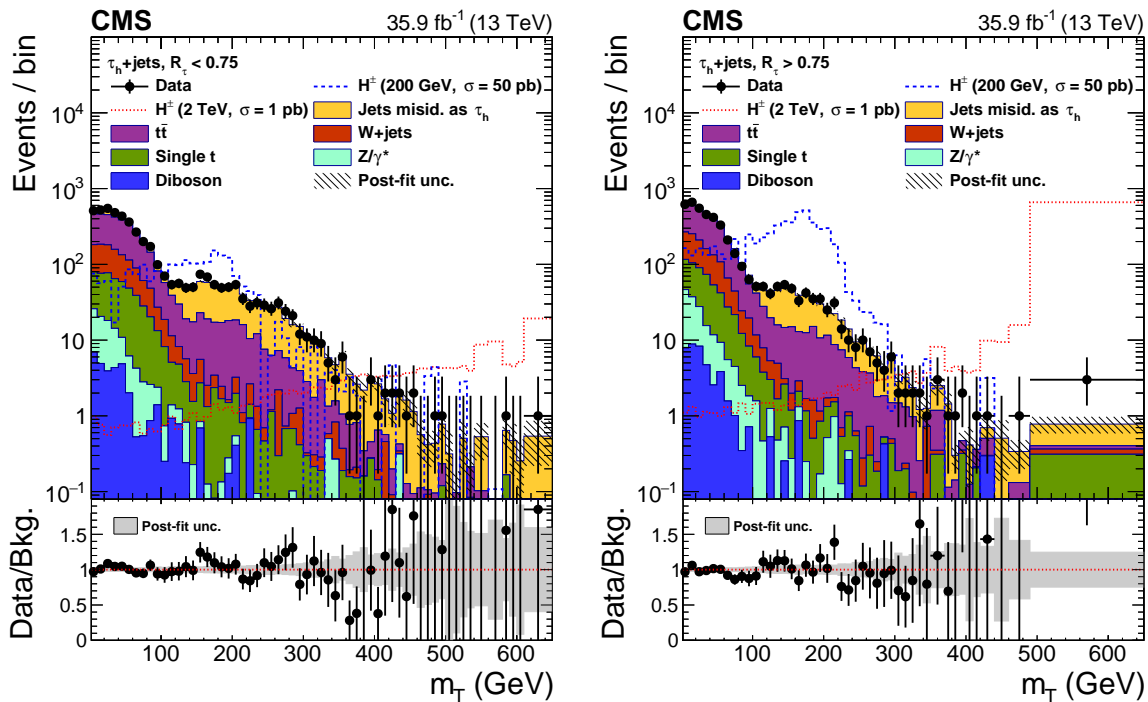
Figure 14.1: The transverse mass distributions in the τ_h + jets final state after a background-only fit to the data. Left: category defined by $R_\tau < 0.75$. Transverse mass values up to 5 TeV are considered in the fit, but the last bins with $m_T > 650\,\text{GeV}$ do not contain any observed events. Right: category defined by $R_\tau > 0.75$. The last bin shown extends to 5 TeV. Two signal samples are overlaid for illustration. [73]

The distributions of m_T after the background-only fit to the data are shown in Figure 14.1 for both R_τ categories. The $t\bar{t}$ background is dominant in the low-m_T region, while the jet $\to \tau_h$ events (mostly QCD multijet events) constitute most of the background in the tail of the distribution. No significant excess is observed, and the data are found to agree with the SM prediction. Two signal samples with $m_{H^\pm}$ of 200 GeV and 2 TeV, normalized to arbitrary cross sections, are overlaid for illustration.

14.2 Exclusion limits

The modified frequentist CL_s criterion [178, 179] based on the profile likelihood ratio test statistic [177] is applied to determine the 95% confidence level (C.L.) limit for the product of the $H^\pm$ production cross section $\sigma_{H^\pm}$ and the branching fraction $\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$, using the reconstructed m_T as the summary statistic. The m_T distributions in the two R_τ categories are fitted simultaneously.

Figure 14.2: The 95% C.L. exclusion limits obtained from the analysis of the τ_h + jets final state for the $H^\pm$ mass range from 80 GeV to 3 TeV. The observed limit on $\sigma_{H^\pm}\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ (solid black points) is compared to the expected limit assuming only Standard Model processes (dashed line). The green (yellow) error bands represent one (two) standard deviations from the expected limit. The horizontal axis is linear from 80 to 180 GeV and logarithmic for larger $m_{H^\pm}$ values.

The asymptotic approximation [181] is used in the limit calculation. The validity of this approximation is confirmed by calculating the limits by generating pseudoexperiments in a subset of mass points.

For the $H^\pm$ mass range up to 165 GeV, the limit is set on $\mathcal{B}(t \to bH^\pm)\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$, scaling down the $t\bar{t}$ background component consistently with the $\mathcal{B}(t \to bH^\pm)$ signal hypothesis. The result is interpreted as a limit on $\sigma_{H^\pm}\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ by assuming $\sigma_{H^\pm} = 2\sigma_{t\bar{t}}\mathcal{B}(t \to bH^\pm)(1 - \mathcal{B}(t \to bH^\pm))$. For the $H^\pm$ mass range from 170 GeV to 3 TeV, the limit on $\sigma_{H^\pm}\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ is calculated without assuming a specific production mode.

The model-independent upper limit with both R_τ categories combined is shown in Figure 14.2. The observed limit ranges from 11 pb at 80 GeV to 4.7 fb at 3 TeV. For the light charged Higgs boson mass range of 80–160 GeV, the limit corresponds to $\mathcal{B}(t \to bH^\pm)\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ values between 0.96% (at 80 GeV) and 0.22% (at 160 GeV).

Figure 14.3: The 95% C.L. exclusion limits obtained from the combined analysis of the hadronic and leptonic final states, for the $H^{\pm}$ mass range from 80 GeV to 3 TeV. The observed limit on $\sigma_{H^{\pm}} \mathcal{B}(H^{\pm} \to \tau^{\pm} \nu_{\tau})$ (solid black points) is compared to the expected limit assuming only Standard Model processes (dashed line). The green (yellow) error bands represent one (two) standard deviations from the expected limit. The horizontal axis is linear from 80 to 180 GeV and logarithmic for larger $m_{H^{\pm}}$ values. [73]

The drop in the expected and observed limits in the intermediate region is not predicted from theory [70], but rather an experimental feature explained by the fact that in this region LO signal samples are used instead of NLO. The dip is mitigated but not completely canceled by the LO-to-NLO corrections extrapolated from the surrounding mass regions, as described in Section 13.7.

14.3 Combination with the leptonic final state

The final results on the $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ decay channel are obtained by combining the results of the analysis in the τ_{h} + jets final state, described in this thesis, with those from the analysis of the leptonic final state, as discussed in Section 10.4.

The final combination is carried out by calculating the model-independent upper limits with the same method as before, but this time also the m_{T} distributions from the

Figure 14.4: Median expected 95% C.L. exclusion limits on $\sigma_{H^\pm}\, \mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ in the hadronic final state (blue), in the leptonic final states (red) and for all final states combined (black). The horizontal axis is linear from 80 to 180 GeV and logarithmic for larger $m_{H^\pm}$ values.

leptonic final states are included in the simultaneous fit. The combined results include the two R_τ categories of the τ_h + jets final state and 34 categories of the leptonic final states. All sources of systematic uncertainties common to the hadronic and leptonic final state analyses are correlated in the fit.

The final model-independent upper limit with all final states and categories combined is shown in Figure 14.3. The numerical values are listed in Table 14.1. The observed limit ranges from 6.0 pb at 80 GeV to 4.8 fb at 3 TeV. For the light charged Higgs boson mass range of 80–160 GeV, the limit corresponds to $\mathcal{B}(t \to bH^\pm)\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ values between 0.36% (at 80 GeV) and 0.08% (at 160 GeV), as detailed in Table 14.2. In the light $H^\pm$ mass range, this is the most stringent limit on $\mathcal{B}(t \to bH^\pm)\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ to date set by the CMS Collaboration, with a factor of 1.5–3.0 improvement with respect to Ref. [27], depending on $m_{H^\pm}$. In the intermediate mass range of 165–175 GeV, this is the first limit on $\sigma_{H^\pm}\mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ set by the CMS Collaboration. In the heavy mass range from 180 GeV, this result extends the search region up to $m_{H^\pm} = 3$ TeV, compared to 600 GeV in Ref. [27].

Table 14.1: The expected and observed 95% C.L. exclusion limits on $\sigma_{H^\pm} \mathcal{B}(H^\pm \to \tau^\pm \nu_\tau)$ for the $H^\pm$ mass range from 80 GeV to 3 TeV. The ± 1 s.d. (± 2 s.d.) refers to one (two) standard deviations from the expected limit.

| $m_{H^\pm}$	Expected limit (pb)					Observed
(GeV)	-2 s.d.	-1 s.d.	median	$+1$ s.d.	$+2$ s.d.	limit (pb)
80	3.17	4.25	5.87	8.15	10.89	5.97
90	3.05	4.08	5.69	7.96	10.75	4.59
100	2.67	3.56	4.94	6.90	9.26	3.24
120	2.04	2.72	3.78	5.29	7.12	2.55
140	1.41	1.87	2.61	3.63	4.88	2.22
150	1.19	1.58	2.20	3.07	4.14	1.63
155	1.06	1.41	1.95	2.71	3.64	1.48
160	1.05	1.39	1.93	2.69	3.61	1.31
165	0.76	1.02	1.45	2.67	2.86	1.01
170	0.40	0.54	0.77	1.12	1.59	0.57
175	0.37	0.50	0.71	1.03	1.45	0.52
180	0.44	0.60	0.83	1.18	1.59	0.85
200	0.30	0.41	0.57	0.80	1.09	0.65
220	0.22	0.30	0.41	0.58	0.80	0.47
250	0.15	0.21	0.29	0.41	0.56	0.31
300	0.08	0.11	0.15	0.22	0.30	0.14
400	0.032	0.043	0.062	0.090	0.125	0.078
500	0.016	0.022	0.031	0.046	0.067	0.048
750	0.0035	0.0050	0.0077	0.012	0.019	0.014
800	0.0029	0.0041	0.0064	0.0102	0.0157	0.0107
1000	0.0020	0.0030	0.0047	0.0077	0.0121	0.0085
2000	0.0009	0.0014	0.0025	0.0044	0.0074	0.0050
2500	0.0007	0.0012	0.0022	0.0042	0.0068	0.0047
3000	0.0007	0.0012	0.0022	0.0043	0.0067	0.0048

Table 14.2: The expected and observed 95% C.L. exclusion limits on $\mathcal{B}(t\rightarrow bH^{\pm})\mathcal{B}(H^{\pm}\rightarrow\tau^{\pm}\nu_{\tau})$ for the $H^{\pm}$ mass range from 80 to 160 GeV. The ± 1 s.d. (± 2 s.d.) refers to one (two) standard deviations from the expected limit.

$m_{H^{\pm}}$	Expected limit (%)					Observed
(GeV)	-2 s.d.	-1 s.d.	median	$+1$ s.d.	$+2$ s.d.	limit (pb)
80	0.19	0.26	0.35	0.49	0.65	0.36
90	0.18	0.25	0.34	0.48	0.65	0.28
100	0.16	0.21	0.30	0.28	0.55	0.20
120	0.12	0.16	0.23	0.31	0.43	0.15
140	0.084	0.112	0.157	0.219	0.294	0.134
150	0.071	0.095	0.132	0.185	0.249	0.098
155	0.063	0.084	0.117	0.163	0.220	0.089
160	0.063	0.084	0.116	0.162	0.217	0.079

Figure 14.4 shows a comparison of the observed limits based on the hadronic final state (blue curve), leptonic final state (red curve) and their combination (black curve). In the light and intermediate $H^{\pm}$ mass regions both hadronic and leptonic final states contribute significantly to the sensitivity, and the combined limits are on average $\approx 40\%$ lower (better) compared to the τ_h + jets final state alone. In the heavy mass region, the sensitivity of the leptonic final states decreases and the τ_h + jets final state starts to dominate the limit as $m_{H^{\pm}}$ increases. The improvement arising from the leptonic final states decreases from 30% at $m_{H^{\pm}} = 180$ GeV to only 3% $m_{H^{\pm}} = 400$ GeV, and above 500 GeV the combined limits match with those from the τ_h + jets final state alone. This behavior is expected based on the different trigger thresholds and jet selection criteria discussed in Section 10.4.

14.4 Interpretation of results in the MSSM

The limits are interpreted in the MSSM $m_h^{\text{mod-}}$ benchmark scenario [63] by comparing the observed and expected limits on $\sigma_{H^{\pm}}\mathcal{B}(H^{\pm}\rightarrow\tau^{\pm}\nu_{\tau})$ to the theoretical cross sections and branching fractions predicted in this scenario [69,70,210–213]. The MSSM $m_h^{\text{mod-}}$ scenario is specified using low-energy MSSM parameters and it is designed to give a mass of approximately 125 GeV for the light CP-even Higgs boson over a wide region of the parameter space. The limit for the MSSM $m_h^{\text{mod-}}$ scenario in the ($m_{H^{\pm}}$, $\tan\beta$) plane is shown in Figure 14.5.

Figure 14.5: The observed 95% C.L. exclusion limits (solid black points), and the expected median (dashed line) in the MSSM $m_h^{\text{mod-}}$ benchmark scenario. The green (yellow) bands represent one (two) standard deviations from the expected limit. The region below the red line is not allowed if we assume that the observed neutral Higgs boson is the light CP-even 2HDM Higgs boson with a mass of $125 \pm 3\,\text{GeV}$, where the uncertainty is the theoretical uncertainty in the mass calculation. [73]

Based on the observed limit, all $\tan\beta$ values from 1 to 60 are excluded for $m_{H^\pm}$ values up to $160\,\text{GeV}$. For $m_{H^\pm} = 200(400)\,\text{GeV}$, the observed limit excludes all $\tan\beta$ values above 26 (40), compared to 45 (56) in Ref. [27].

In Figure 14.5, the red line indicates the limits for the possible $m_{H^\pm}$ and $\tan\beta$ values set by the observation of the $125\,\text{GeV}$ Higgs boson. Even though this scenario is designed to be compatible with its measured mass and couplings in a wide region of the parameter space [63], the areas below and left of the red line are excluded assuming that the measured Higgs boson is the lightest CP-odd neutral scalar boson of the MSSM.

In Figure 14.6, the exclusion limit in the MSSM $m_h^{\text{mod-}}$ scenario is compared to a recent CMS analysis targeting the $H^\pm \to tb$ search channel [74]. In terms of $H^\pm$ branching fraction, in this scenario these two decay channels are among the dominant ones as discussed in Section 3.2.2.

Figure 14.6: The observed 95% C.L. exclusion limits (solid points), and the expected median (dashed line) from the $H^{\pm} \to \tau^{\pm}\nu_{\tau}$ decay channel presented in this thesis (blue) and from a recent search targeting the $H^{\pm} \to tb$ channel (black), compared in the MSSM $m_h^{\mathrm{mod-}}$ benchmark scenario. The region below the red line is not allowed if we assume that the observed neutral Higgs boson is the light CP-even 2HDM Higgs boson with a mass of $125 \pm 3\,\mathrm{GeV}$, where the uncertainty is the theoretical uncertainty in the mass calculation. [74]

Several interesting features of the $\tau^{\pm}\nu_{\tau}$ and tb channels can be seen in Figure 14.6. Firstly, these search channels complement each other, since the limit on $H^{\pm} \to tb$ constrains both very small and very large $\tan\beta$ values, whereas the limit on $H^{\pm} \to \tau^{\pm}\nu_{\tau}$ constrains only the latter. In this particular scenario, however, the indirect limit from the observed 125 GeV boson is stronger in the low-$\tan\beta$ region than the limit obtained from the tb decay channel. Secondly, even in the large-$m_{H^{\pm}}$ region where the $H^{\pm}$ branching fraction to tb is significantly larger than the branching fraction to $\tau^{\pm}\nu_{\tau}$, the results from the $\tau^{\pm}\nu_{\tau}$ channel set a more stringent limit for large $\tan\beta$ values.

The competitiveness of the subdominant $\tau^{\pm}\nu_{\tau}$ channel is largely due to good discrimination power provided by the transverse mass distribution of the $\tau\nu$ system, combined with additional handles such as angular selections and the R_{τ} variable reflecting the tau lepton helicity.

As a final remark, we note that a significant part of the $(m_{H^\pm}, \tan\beta)$ plane, corresponding to high charged Higgs boson masses and large $\tan\beta$ values, remains compatible with these recent search CMS results as well as the previous experimental constraints discussed in Section 3.2.3, leaving space for future analysis efforts with larger data samples and continuously improving analysis methods.

Chapter 15

Summary and outlook

Extensions of the Standard Model with at least two Higgs doublets predict the existence of electrically charged Higgs bosons. The observation of charged Higgs bosons would provide direct evidence for physics beyond the Standard Model, and guide our way towards a wider theory. On the other hand, when results of an experimental search agree with the Standard Model, they can be used to constrain or exclude proposed theoretical models with extended Higgs sectors.

In this thesis, a search is presented for charged Higgs bosons decaying into a tau lepton and a neutrino, based on proton-proton collision events recorded by the CMS experiment in 2016 at a center-of-mass energy of 13 TeV. In many models, the $H^\pm \to \tau^\pm \nu_\tau$ channel is considered as one of the most sensitive for the $H^\pm$ discovery. The search targets the hadronic final state with a hadronically decaying tau lepton, missing transverse momentum due to neutrinos, and additional jets from top quark decays.

The data used in the search correspond to an integrated luminosity of $35.9\,\text{fb}^{-1}$, almost twice as much as the 8 TeV data used in the previous CMS analysis on this channel. The search covers a wide range of $H^\pm$ mass hypotheses from 80 GeV up to 3 TeV. For the first time in CMS, the intermediate mass range with $H^\pm$ mass values close to the top quark mass is also included in the search. The signal is modeled at next-to-leading order in perturbative quantum chromodynamics, except for the intermediate mass range where the current signal generation model is limited to leading order. In the intermediate mass range, a correction is derived to compensate for the difference in acceptance between the leading and next-to-leading order signal samples.

A dedicated $\tau_\text{h} + p_\text{T}^\text{miss}$ trigger is used for the online event selection. The particle-identification algorithms and selection criteria applied offline are carefully chosen for good performance under challenging pileup conditions. In the offline selection, the

events are required to contain an isolated τ_h candidate, high p_T^{miss}, and at least three hadronic jets. At least one of the jets is required to pass a b jet identification. Angular cuts based on the azimuthal angle differences between the τ_h candidate, p_T^{miss} and the leading jets are applied to reduce the QCD multijet background. A variable reflecting the helicity of the tau leptons is used to categorize the events in the statistical analysis.

The background from events with jets misidentified as tau leptons is estimated from data, whereas the background from events with genuine tau leptons is estimated from simulation. The jet $\to \tau_h$ background is dominated by the QCD multijet events, whereas the genuine-tau background mostly contains $t\bar{t}$ events with a $W^{\pm}$ boson decaying into a tau lepton. A new version of the tau embedding method, which can be used to estimate the genuine-tau background from data using single-muon events, is also presented. A third, small background contribution arises from events with electrons or muons misidentified as τ_h, and is estimated from simulation.

The transverse mass distribution of the tau-neutrino system is reconstructed and used to test the compatibility of the observed distribution against the signal and background hypotheses. The results agree with the background-only expectation from the Standard Model, so the transverse mass distributions are used to set upper limits for the product of the $H^{\pm}$ production cross section and the branching fraction to $\tau^{\pm}\nu_{\tau}$ at 95% confidence level. The modified frequentist CL_s criterion based on the profile likelihood ratio test statistic is used. For maximal signal sensitivity, the results of this analysis are combined with those from the leptonic final states of the $H^{\pm} \to \tau^{\pm}\nu_{\tau}$ channel. The systematic uncertainties are incorporated as nuisance parameters in the likelihood. For the combined result, the observed limit ranges from 6.0 pb at 80 GeV to 4.8 fb at 3 TeV.

The results can be interpreted in chosen theoretical models by comparing the observed limit to the cross sections and branching fractions predicted by the corresponding model. In this thesis, the results are interpreted as constraints in the parameter space of the Minimal Supersymmetric Standard Model m_h^{mod-} benchmark scenario. In this scenario, the observed limit excludes all $\tan\beta$ values from 1 to 60 for charged Higgs boson masses up to 160 GeV, and above 160 GeV it sets more stringent restrictions for the allowed $\tan\beta$ range than any previous CMS result in the $H^{\pm} \to \tau^{\pm}\nu_{\tau}$ channel.

This first analysis of the $H^{\pm} \to \tau^{\pm}\nu_{\tau}$ channel based on 13 TeV collision data sets the ground for further searches which will employ the full 13 TeV data set of approximately 140 fb^{-1} recorded at the CMS in 2016–2018. While four times larger data set will

double the statistical precision, its impact on the analysis of the $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ channel is not expected to be that large, since the sensitivity of this channel is also limited by systematic uncertainties. Thus methodological improvements remain essential in pushing the boundaries of the sensitivity further.

The efforts to improve analysis methods will continue in three fronts. Firstly, modern multivariate analysis techniques, such as deep neural networks, can be used to improve the discrimination between the background and the signal events, improving the overall signal-to-background ratio. Secondly, the statistical analysis can be further improved by providing the maximum-likelihood fit with more information as input, for example by classifying the events into several categories. Inclusion of additional signal-depleted categories can be useful to constrain the dominating systematic uncertainties. Thirdly, improved background estimation techniques such as the tau embedding method may provide more accurate background estimates with smaller associated systematic uncertainties compared to the current methods. Just as the work presented in this thesis is built upon the previous analysis efforts, hopefully the methods and results presented here will provide stepping stones towards the next improvements.

The ambitious LHC physics program will continue for decades. After the end of Run 3, planned for 2021–2023, the amount of collected data is expected to exceed $300 \, \mathrm{fb}^{-1}$. Then the next major milestone will be the installation of the high-luminosity LHC (HL-LHC), which is expected to deliver $3000 \, \mathrm{fb}^{-1}$ of data by 2035. The efforts to upgrade the CMS detector to cope with the harsh luminosity conditions of the HL-LHC are already ongoing. The upgraded CMS detector will be faster, more radiation-hard, of higher granularity, and equipped with a completely new trigger system.

The continuously increasing amount of data will allow extremely precise measurements of the properties of the known particles such as the $125 \, \mathrm{GeV}$ Higgs boson and the top quark, as well as ambitious searches for new physics, including charged Higgs bosons. The exciting thing about high-energy physics is that as we probe nature with an unprecedented precision, even the best theories can only provide educated guesses about what will be found—the complicated experimental arrangements and involved statistical analysis are the only way to really find out. If we are lucky, some unexpected discovery will revolutionize our understanding of nature in the years to come.

Bibliography

[1] S. L. Glashow, "Partial symmetries of weak interactions", *Nucl. Phys.* **22** (1961) 579, `doi:10.1016/0029-5582(61)90469-2`.

[2] S. Weinberg, "A model of leptons", *Phys. Rev. Lett.* **19** (1967) 1264, `doi:10.1103/PhysRevLett.19.1264`.

[3] A. Salam, "Weak and electromagnetic interactions", in *Proceedings of the Eigth Nobel Symposium*, p. 367. 1968.

[4] M. Gell-Mann, "The Eightfold Way: A theory of strong interaction symmetry", in *World Scientific Series in 20th Century Physics*. 2010. Originally published in 1961.

[5] G. Zweig, "An SU(3) model for strong interaction symmetry and its breaking. Version 2", in *Developments in the Quark Theory of Hadrons. Volume 1. 1964–1978*, p. 22. Hadronic Press, 1980.

[6] H. Fritzsch, M. Gell-Mann, and H. Leutwyler, "Advantages of the color octet gluon picture", *Phys. Lett.* **47B** (1973) 365, `doi:10.1016/0370-2693(73)90625-4`.

[7] P. W. Higgs, "Broken symmetries, massless particles and gauge fields", *Phys. Lett.* **12** (1964) 132, `doi:10.1016/0031-9163(64)91136-9`.

[8] P. W. Higgs, "Broken symmetries and the masses of gauge bosons", *Phys. Rev. Lett.* **13** (1964) 508, `doi:10.1103/PhysRevLett.13.508`.

[9] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, "Global conservation laws and massless particles", *Phys. Rev. Lett.* **13** (1964) 585, `doi:10.1103/PhysRevLett.13.585`.

[10] P. W. Higgs, "Spontaneous symmetry breakdown without massless bosons", *Phys. Rev.* **145** (1966) 1156, `doi:10.1103/PhysRev.145.1156`.

[11] T. W. B. Kibble, "Symmetry breaking in non-Abelian gauge theories", *Phys. Rev.* **155** (1967) 1554, `doi:10.1103/PhysRev.155.1554`.

[12] F. Englert and R. Brout, "Broken symmetry and the mass of gauge vector mesons", *Phys. Rev. Lett.* **13** (1964) 321, `doi:10.1103/PhysRevLett.13.321`.

[13] L. Evans and P. Bryant, "LHC machine", *JINST* **3** (2008) S08001, `doi:10.1088/1748-0221/3/08/S08001`.

[14] ATLAS Collaboration, "Observation of a new particle in the search for the standard model Higgs boson with the ATLAS detector at the LHC", *Phys. Lett. B* **716** (2012) 1, `doi:10.1016/j.physletb.2012.08.020`, `arXiv:1207.7214`.

[15] CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC", *Phys. Lett. B* **716** (2012) 30, `doi:10.1016/j.physletb.2012.08.021`, `arXiv:1207.7235`.

[16] CMS Collaboration, "Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **06** (2013) 081, `doi:10.1007/JHEP06(2013)081`, `arXiv:1303.4571`.

[17] T. D. Lee, "A theory of spontaneous T violation", *Phys. Rev. D* **8** (1973) 1226, `doi:10.1103/PhysRevD.8.1226`.

[18] G. C. Branco et al., "Theory and phenomenology of two-Higgs-doublet models", *Phys. Rept.* **516** (2012) 1, `doi:10.1016/j.physrep.2012.02.002`, `arXiv:1106.0034`.

[19] ALEPH, DELPHI, L3, OPAL and LEP Collaborations, "Search for charged Higgs bosons: combined results using LEP data", *Eur. Phys. J. C* **73** (2013) 2463, `doi:10.1140/epjc/s10052-013-2463-1`, `arXiv:1301.6065`.

[20] CDF Collaboration, "Search for Higgs bosons predicted in two-Higgs-doublet models via decays to tau lepton pairs in 1.96 TeV $p\bar{p}$ collisions", *Phys. Rev. Lett.* **103** (2009) 201801, `doi:10.1103/PhysRevLett.103.201801`, `arXiv:0906.1014`.

[21] D0 Collaboration, "Search for Higgs bosons of the minimal supersymmetric standard model in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV", *Phys. Lett. B* **710** (2012) 569, `doi:10.1016/j.physletb.2012.03.021`, `arXiv:1112.5431`.

[22] ATLAS Collaboration, "Search for charged Higgs bosons decaying via $H^+ \to \tau\nu$ in top quark pair events using pp collision data at $\sqrt{s} = 7$ TeV with the ATLAS detector", *JHEP* **06** (2012) 039, `doi:10.1007/JHEP06(2012)039`, `arXiv:1204.2760`.

[23] CMS Collaboration, "Search for a light charged Higgs boson in top quark decays in pp collisions at $\sqrt{s} = 7$ TeV", *JHEP* **07** (2012) 143, `doi:10.1007/JHEP07(2012)143`, `arXiv:1205.5736`.

[24] ATLAS Collaboration, "Search for charged Higgs bosons through the violation of lepton universality in $t\bar{t}$ events using pp collision data at $\sqrt{s} = 7$ TeV with the ATLAS experiment", *JHEP* **03** (2013) 076, `doi:10.1007/JHEP03(2013)076`, `arXiv:1212.3572`.

[25] ATLAS Collaboration, "Search for a light charged Higgs boson in the decay channel $H^+ \to c\bar{s}$ in $t\bar{t}$ events using pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector", *Eur. Phys. J. C* **73** (2013) 2465, `doi:10.1140/epjc/s10052-013-2465-z`, `arXiv:1302.3694`.

[26] ATLAS Collaboration, "Search for charged Higgs bosons decaying via $H^\pm \to \tau^\pm \nu$ in fully hadronic final states using pp collision data at $\sqrt{s} = 8$ TeV with the ATLAS detector", *JHEP* **03** (2015) 088, `doi:10.1007/JHEP03(2015)088`, `arXiv:1412.6663`.

[27] CMS Collaboration, "Search for a charged Higgs boson in pp collisions at $\sqrt{s} = 8$ TeV", *JHEP* **11** (2015) 018, `doi:10.1007/JHEP11(2015)018`, `arXiv:1508.07774`.

[28] ATLAS Collaboration, "Search for charged Higgs bosons in the $H^\pm \to tb$ decay channel in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector", *JHEP* **03** (2016) 127, `doi:10.1007/JHEP03(2016)127`, `arXiv:1512.03704`.

[29] CMS Collaboration, "Search for a light charged Higgs boson decaying to $c\bar{s}$ in pp collisions at $\sqrt{s} = 8$ TeV", *JHEP* **12** (2015) 178, `doi:10.1007/JHEP12(2015)178`, `arXiv:1510.04252`.

[30] ATLAS Collaboration, "Search for charged Higgs bosons produced in association with a top quark and decaying via $H^\pm \to \tau\nu$ using pp collision data recorded at $\sqrt{s} = 13$ TeV by the ATLAS detector", *Phys. Lett. B* **759** (2016) 555, `doi:10.1016/j.physletb.2016.06.017`, `arXiv:1603.09203`.

[31] ATLAS Collaboration, "Search for charged Higgs bosons decaying via $H^{\pm} \to \tau^{\pm} \nu_{\tau}$ in the τ+jets and τ+lepton final states with 36 fb^{-1} of pp collision data recorded at $\sqrt{s} = 13$ TeV with the ATLAS experiment", *JHEP* **09** (2018) 139, `doi:10.1007/JHEP09(2018)139`, `arXiv:1807.07915`.

[32] ATLAS Collaboration, "Search for charged Higgs bosons decaying into top and bottom quarks at $\sqrt{s} = 13$ TeV with the ATLAS detector", *JHEP* **11** (2018) 085, `doi:10.1007/JHEP11(2018)085`, `arXiv:1808.03599`.

[33] CMS Collaboration, "Search for a charged Higgs boson decaying to charm and bottom quarks in proton-proton collisions at $\sqrt{s} = 8$ TeV", *JHEP* **11** (2018) 115, `doi:10.1007/JHEP11(2018)115`, `arXiv:1808.06575`.

[34] A. Djouadi, "The anatomy of electro-weak symmetry breaking. II The Higgs bosons in the minimal supersymmetric model", *Phys. Rept.* **459** (2008) 1, `doi:10.1016/j.physrep.2007.10.005`, `arXiv:hep-ph/0503173`.

[35] Particle Data Group Collaboration, "Review of particle physics", *Phys. Rev. D* **98** (2018) 030001, `doi:10.1103/PhysRevD.98.030001`.

[36] I. J. R. Aitchison and A. J. G. Hey, "Gauge theories in particle physics: A practical introduction. Volume 2: Non-Abelian gauge theories: QCD and the electroweak theory". CRC Press, Bristol, UK, 2012.

[37] M. D. Schwartz, "Quantum field theory and the Standard Model". Cambridge University Press, 2014.

[38] W. Pauli, "The connection between spin and statistics", *Phys. Rev.* **58** (1940) 716, `doi:10.1103/PhysRev.58.716`.

[39] CMS Collaboration, "Constraints on the spin-parity and anomalous HVV couplings of the Higgs boson in proton collisions at 7 and 8 TeV", *Phys. Rev. D* **92** (2015) 012004, `doi:10.1103/PhysRevD.92.012004`, `arXiv:1411.3441`.

[40] ATLAS and CMS Collaborations, "Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments", *Phys. Rev. Lett.* **114** (2015) 191803, `doi:10.1103/PhysRevLett.114.191803`, `arXiv:1503.07589`.

[41] ATLAS Collaboration, "Study of the spin and parity of the Higgs boson in diboson decays with the ATLAS detector", *Eur. Phys. J. C* **75** (2015) 476, `doi:10.1140/epjc/s10052-015-3685-1`, `arXiv:1506.05669`. [Erratum: *Eur. Phys. J. C* **75** (2015) 476].

[42] ATLAS and CMS Collaborations, "Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **08** (2016) 045, `doi:10.1007/JHEP08(2016)045`, `arXiv:1606.02266`.

[43] CMS Collaboration, "Measurements of properties of the Higgs boson decaying into the four-lepton final state in pp collisions at $\sqrt{s} = 13$ TeV", *JHEP* **11** (2017) 047, `doi:10.1007/JHEP11(2017)047`, `arXiv:1706.09936`.

[44] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak Interactions with lepton-hadron symmetry", in *Meeting of the Italian School of Physics and Weak Interactions Bologna, Italy, April 26–28, 1984*, volume 2, p. 1285. 1970. `doi:10.1103/PhysRevD.2.1285`.

[45] S. Raychaudhuri and D. P. Roy, "Sharpening up the charged Higgs boson signature using τ polarization at LHC", *Phys. Rev. D* **53** (1996) 4902, `doi:10.1103/PhysRevD.53.4902`, `arXiv:hep-ph/9507388`.

[46] D. P. Roy, "The Hadronic τ decay signature of a heavy charged Higgs boson at LHC", *Phys. Lett. B* **459** (1999) 607, `doi:10.1016/S0370-2693(99)00724-8`, `arXiv:hep-ph/9905542`.

[47] D. P. Roy, "Looking for the charged Higgs boson", *Mod. Phys. Lett. A* **19** (2004) 1813, `doi:10.1142/S0217732304015105`, `arXiv:hep-ph/0406102`.

[48] M. Guchait, R. Kinnunen, and D. P. Roy, "Signature of heavy charged Higgs boson at LHC in the 1 and 3 prong hadronic tau decay channels", *Eur. Phys. J. C* **52** (2007) 665, `doi:10.1140/epjc/s10052-007-0396-2`, `arXiv:hep-ph/0608324`.

[49] Y. Nagashima, "Beyond the standard model of elementary particle physics". Wiley-VCH, Weinheim, 2014.

[50] A. D. Sakharov, "Violation of CP onvariance, C asymmetry, and baryon asymmetry of the universe", *Pisma Zh. Eksp. Teor. Fiz.* **5** (1967) 32, `doi:10.1070/PU1991v034n05ABEH002497`.

[51] W.-L. Guo and Y.-L. Wu, "The real singlet scalar dark matter model", *JHEP* **10** (2010) 083, `doi:10.1007/JHEP10(2010)083`, `arXiv:1006.2518`.

[52] S. Dawson and M. Sullivan, "Enhanced di-Higgs boson production in the complex Higgs singlet model", *Phys. Rev. D* **97** (2018), no. 1, 015022, `doi:10.1103/PhysRevD.97.015022`, `arXiv:1711.06683`.

[53] A. Drozd, B. Grzadkowski, and J. Wudka, "Multi-scalar-singlet extension of the Standard Model: the case for dark matter and an invisible Higgs boson", *JHEP* **04** (2012) 006, `doi:10.1007/JHEP04(2012)006,10.1007/JHEP11(2014)130`, `arXiv:1112.2582`. [Erratum: *JHEP* **11** (2014) 130].

[54] G. C. Dorsch, S. J. Huber, and J. M. No, "A strong electroweak phase transition in the 2HDM after LHC8", *JHEP* **10** (2013) 029, `doi:10.1007/JHEP10(2013)029`, `arXiv:1305.6610`.

[55] J. O. Andersen et al., "Nonperturbative analysis of the electroweak phase transition in the two Higgs doublet model", *Phys. Rev. Lett.* **121** (2018), no. 19, 191802, `doi:10.1103/PhysRevLett.121.191802`, `arXiv:1711.09849`.

[56] N. Blinov, S. Profumo, and T. Stefaniak, "The electroweak phase transition in the Inert Doublet Model", *JCAP* **1507** (2015), no. 07, 028, `doi:10.1088/1475-7516/2015/07/028`, `arXiv:1504.05949`.

[57] J. F. Gunion et al., "The Higgs Hunter's Guide". Brookhaven Nat. Lab., 1989.

[58] S. L. Glashow and S. Weinberg, "Natural conservation laws for neutral currents", *Phys. Rev. D* **15** (1977) 1958, `doi:10.1103/PhysRevD.15.1958`.

[59] I. F. Ginzburg, M. Krawczyk, and P. Osland, "Two Higgs doublet models with CP violation", in *Linear colliders. Proceedings, International Workshop on physics and experiments with future electron-positron linear colliders, LCWS 2002, Seogwipo, Jeju Island, Korea, August 26–30, 2002*, p. 703. 2002. `arXiv:hep-ph/0211371`.

[60] V. D. Barger, J. L. Hewett, and R. J. N. Phillips, "New constraints on the charged Higgs sector in two Higgs doublet models", *Phys. Rev. D* **41** (1990) 3421, `doi:10.1103/PhysRevD.41.3421`.

[61] G. C. Branco, W. Grimus, and L. Lavoura, "Relating the scalar flavor changing neutral couplings to the CKM matrix", *Phys. Lett. B* **380** (1996) 119, `doi:10.1016/0370-2693(96)00494-7`, `arXiv:hep-ph/9601383`.

[62] S. Weinberg, "The Quantum Theory of Fields. Volume 3: Supersymmetry". Cambridge University Press, 2013.

[63] M. Carena et al., "MSSM Higgs boson searches at the LHC: Benchmark scenarios after the discovery of a Higgs-like particle", *Eur. Phys. J. C* **73** (2013) 2552, `doi:10.1140/epjc/s10052-013-2552-1`, `arXiv:1302.7033`.

[64] LHC Higgs Cross Section Working Group Collaboration, "Handbook of LHC Higgs Cross Sections: 3. Higgs Properties", `doi:10.5170/CERN-2013-004`, `arXiv:1307.1347`.

[65] A. G. Akeroyd et al., "Prospects for charged Higgs searches at the LHC", *Eur. Phys. J. C* **77** (2017), no. 5, 276, `doi:10.1140/epjc/s10052-017-4829-2`, `arXiv:1607.01320`.

[66] G. Senjanovic and R. N. Mohapatra, "Exact left-right symmetry and spontaneous violation of parity", *Phys. Rev. D* **12** (1975) 1502, `doi:10.1103/PhysRevD.12.1502`.

[67] J. F. Gunion, R. Vega, and J. Wudka, "Higgs triplets in the standard model", *Phys. Rev. D* **42** (1990) 1673, `doi:10.1103/PhysRevD.42.1673`.

[68] H. Georgi and M. Machacek, "Doubly charged Higgs bosons", *Nucl. Phys. B* **262** (1985) 463, `doi:10.1016/0550-3213(85)90325-6`.

[69] LHC Higgs Cross Section Working Group Collaboration, "Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector", `arXiv:1610.07922`.

[70] C. Degrande et al., "Accurate predictions for charged Higgs production: Closing the $m_{\mathrm{H}^\pm} \sim m_{\mathrm{t}}$ window", *Phys. Lett. B* **772** (2017) 87, `doi:10.1016/j.physletb.2017.06.037`, `arXiv:1607.05291`.

[71] A. Arbey et al., "Status of the charged Higgs boson in two Higgs doublet models", *Eur. Phys. J. C* **78** (2018), no. 3, 182, `doi:10.1140/epjc/s10052-018-5651-1`, `arXiv:1706.07414`.

[72] CMS Collaboration, "Search for charged Higgs bosons with the $\mathrm{H}^\pm \to \tau^\pm \nu_\tau$ decay channel in the fully hadronic final state at $\sqrt{s} = 13$ TeV", CMS Physics Analysis Summary CMS-PAS-HIG-16-031, 2016.

[73] CMS Collaboration, "Search for charged Higgs bosons in the $H^\pm \to \tau^\pm \nu_\tau$ decay channel in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **07** (2019) 142, `doi:10.1007/JHEP07(2019)142`, `arXiv:1903.04560`.

[74] CMS Collaboration, "Search for a charged Higgs boson decaying into top and bottom quarks in proton-proton collisions at 13 TeV in events with electrons or muons", CMS Physics Analysis Summary CMS-PAS-HIG-18-004, 2019.

[75] CMS Collaboration, "Search for charged Higgs bosons decaying into top and a bottom quark in the fully hadronic final state at 13 TeV", CMS Physics Analysis Summary CMS-PAS-HIG-18-015, 2019.

[76] ATLAS Collaboration, "Search for a charged Higgs boson produced in the vector-boson fusion mode with decay $H^\pm \to W^\pm Z$ using pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS experiment", *Phys. Rev. Lett.* **114** (2015) 231801, `doi:10.1103/PhysRevLett.114.231801`, `arXiv:1503.04233`.

[77] CMS Collaboration, "Search for charged Higgs bosons produced via vector boson fusion and decaying into a pair of W and Z bosons using pp collisions at $\sqrt{s} = 13$ TeV", *Phys. Rev. Lett.* **119** (2017) 141802, `doi:10.1103/PhysRevLett.119.141802`, `arXiv:1705.02942`.

[78] CMS Collaboration, "Measurement of electroweak WZ boson production and search for new physics in WZ + two jets events in pp collisions at $\sqrt{s} = 13$TeV", *Phys. Lett. B* **795** (2019) 281, `doi:10.1016/j.physletb.2019.05.042`, `arXiv:1901.04060`.

[79] CMS Collaboration, "Observation of electroweak production of same-sign W boson pairs in the two jet and two same-sign lepton final state in proton-proton collisions at $\sqrt{s} = 12$ TeV", *Phys. Rev. Lett.* **120** (2018), no. 8, 081801, `doi:10.1103/PhysRevLett.120.081801`, `arXiv:1709.05822`.

[80] ATLAS Collaboration, "Search for doubly charged scalar bosons decaying into same-sign W boson pairs with the ATLAS detector", *Eur. Phys. J. C* **79** (2019), no. 1, 58, `doi:10.1140/epjc/s10052-018-6500-y`, `arXiv:1808.01899`.

[81] CMS Collaboration, "Summary results of high mass BSM Higgs searches using CMS Run-I data", CMS Physics Analysis Summary CMS-PAS-HIG-16-007, 2016.

[82] 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", *JHEP* **01** (2018) 055, `doi:10.1007/JHEP01(2018)055`, `arXiv:1709.07242`.

[83] CMS Collaboration, "Search for additional neutral MSSM Higgs bosons in the $\tau\tau$ final state in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **09** (2018) 007, `doi:10.1007/JHEP09(2018)007`, `arXiv:1803.06553`.

[84] CMS Collaboration, "Search for beyond the standard model Higgs bosons decaying into a b$\bar{\text{b}}$ pair in pp collisions at $\sqrt{s} = 13$ TeV", *JHEP* **08** (2018) 113, `doi:10.1007/JHEP08(2018)113`, `arXiv:1805.12191`.

[85] ATLAS Collaboration, "Search for heavy Higgs bosons A/H decaying to a top quark pair in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector", *Phys. Rev. Lett.* **119** (2017), no. 19, 191803, `doi:10.1103/PhysRevLett.119.191803`, `arXiv:1707.06025`.

[86] ATLAS Collaboration, "Search for the dimuon decay of the Higgs boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", *Phys. Rev. Lett.* **119** (2017), no. 5, 051802, `doi:10.1103/PhysRevLett.119.051802`, `arXiv:1705.04582`.

[87] CMS Collaboration, "Search for MSSM Higgs bosons decaying to $\mu^+\mu^-$ in proton-proton collisions at $\sqrt{s} = 13$ TeV", `arXiv:1907.03152`. Submitted to *Phys. Lett. B*.

[88] CMS Collaboration, "Search for high-mass diphoton resonances in proton–proton collisions at 13 TeV and combination with 8 TeV search", *Phys. Lett. B* **767** (2017) 147, `doi:10.1016/j.physletb.2017.01.027`, `arXiv:1609.02507`.

[89] ATLAS Collaboration, "Search for new phenomena in high-mass diphoton final states using 37 fb^{-1} of proton–proton collisions collected at $\sqrt{s} = 13$ TeV with the ATLAS detector", *Phys. Lett. B* **775** (2017) 105, `doi:10.1016/j.physletb.2017.10.039`, `arXiv:1707.04147`.

[90] CMS Collaboration, "Search for a fermiophobic Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV", *JHEP* **09** (2012) 111, `doi:10.1007/JHEP09(2012)111`, `arXiv:1207.1130`.

[91] ATLAS Collaboration, "Search for a high-mass Higgs boson decaying to a W boson pair in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector", *JHEP* **01** (2016) 032, `doi:10.1007/JHEP01(2016)032`, `arXiv:1509.00389`.

[92] ATLAS Collaboration, "Search for an additional, heavy Higgs boson in the H → ZZ decay channel at $\sqrt{s} = 8$ TeV in pp collision data with the ATLAS detector", *Eur. Phys. J. C* **76** (2016) 45, `doi:10.1140/epjc/s10052-015-3820-z`, `arXiv:1507.05930`.

[93] CMS Collaboration, "Search for new diboson resonances in the dilepton + jets final state at $\sqrt{s} = 13$ TeV with 2016 data", CMS Physics Analysis Summary CMS-PAS-HIG-16-034, 2017.

[94] CMS Collaboration, "Search for neutral resonances decaying into a Z boson and a pair of b jets or τ leptons", *Phys. Lett. B* **759** (2016) 369, `doi:10.1016/j.physletb.2016.05.087`, `arXiv:1603.02991`.

[95] CMS Collaboration, "Search for a heavy pseudoscalar boson decaying to a Z and a Higgs boson at $\sqrt{s} = 13$ TeV", *Eur. Phys. J. C* **79** (2019) 564, `doi:10.1140/epjc/s10052-019-7058-z`, `arXiv:1903.00941`.

[96] CMS Collaboration, "Search for a heavy pseudoscalar Higgs boson decaying into a 125 GeV Higgs boson and a Z boson in final states of two light leptons and two tau leptons at $\sqrt{s} = 13$ TeV", CMS Physics Analysis Summary CMS-PAS-HIG-18-023, 2019.

[97] ATLAS Collaboration, "Search for the Higgs boson produced in association with a W boson and decaying to four b-quarks via two spin-zero particles in pp collisions at 13 TeV with the ATLAS detector", *Eur. Phys. J. C* **76** (2016), no. 11, 605, `doi:10.1140/epjc/s10052-016-4418-9`, `arXiv:1606.08391`.

[98] CMS Collaboration, "Search for light bosons in decays of the 125 GeV Higgs boson in proton-proton collisions at $\sqrt{s} = 8$ TeV", *JHEP* **10** (2017) 076, `doi:10.1007/JHEP10(2017)076`, `arXiv:1701.02032`.

[99] CMS Collaboration, "A search for pair production of new light bosons decaying into muons in proton-proton collisions at 13 TeV", *Phys. Lett. B* **796** (2019) 131–154, `doi:10.1016/j.physletb.2019.07.013`, `arXiv:1812.00380`.

[100] J. Haller et al., "Update of the global electroweak fit and constraints on two-Higgs-doublet models", *Eur. Phys. J. C* **78** (2018), no. 8, 675, `doi:10.1140/epjc/s10052-018-6131-3`, `arXiv:1803.01853`.

[101] M. E. Peskin and T. Takeuchi, "A new constraint on a strongly interacting Higgs sector", *Phys. Rev. Lett.* **65** (1990) 964, `doi:10.1103/PhysRevLett.65.964`.

[102] M. E. Peskin and T. Takeuchi, "Estimation of oblique electroweak corrections", *Phys. Rev. D* **46** (1992) 381, `doi:10.1103/PhysRevD.46.381`.

[103] J.-H. Park, "Lepton non-universality at LEP and charged Higgs", *JHEP* **10** (2006) 077, `doi:10.1088/1126-6708/2006/10/077`, `arXiv:hep-ph/0607280`.

[104] W. Mader et al., "LHC explores what LEP hinted at: CP-violating Type-I 2HDM", *JHEP* **09** (2012) 125, `doi:10.1007/JHEP01(2014)006,10.1007/JHEP09(2012)125`, `arXiv:1205.2692`. [Erratum: *JHEP* **01** (2014) 006].

[105] W. Stirling, "Parton luminosity and cross section plots". Online at `http://www.hep.ph.ic.ac.uk/~wstirlin/plots/plots.html`, 2012. Accessed 29 May 2019.

[106] J.-P. Tock et al., "The second LHC long shutdown (LS2) for the superconducting magnets", in *Proceedings, 9th International Particle Accelerator Conference (IPAC 2018): Vancouver, BC Canada, April 29–May 4, 2018*. 2018. `doi:10.18429/JACoW-IPAC2018-MOPMF056`.

[107] E. Mobs, "The CERN accelerator complex",.

[108] P. Gagnon, "Full moon pulls LHC from its protons". Online at `https://home.cern/news/news/accelerators/full-moon-pulls-lhc-its-protons`, 2008. Accessed 24 July 2019.

[109] CMS Collaboration, "CMS Luminosity – Public Results". Online at `https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults`, 2019. Accessed 25 April 2019.

[110] CMS Collaboration, "CMS detector design illustration". Online at `http://cms.web.cern.ch/news/cms-detector-design`, 2019. Accessed 26 May 2019.

[111] CMS Collaboration, "The CMS experiment at the CERN LHC", *JINST* **3** (2008) S08004, `doi:10.1088/1748-0221/3/08/S08004`.

[112] A. Tricomi, "Upgrade of the CMS tracker", *JINST* **9** (2014) C03041, `doi:10.1088/1748-0221/9/03/C03041`.

[113] CMS Collaboration, "The performance of the CMS muon system with data at $\sqrt{s}$ = 13 TeV", Technical Report CMS-CR-2018-167, 2018.

[114] CMS Collaboration, "CMS Technical Design Report for the Level-1 Trigger Upgrade", CMS Technical Design Report CERN-LHCC-2013-011, CMS-TDR-012, 2013.

[115] CMS Collaboration, "CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter", Technical Report CERN-LHCC-2012-015, CMS-TDR-010.

[116] CMS Collaboration, "CMS Technical Design Report for the Pixel Detector Upgrade", technical report, 2012.

[117] CMS Collaboration, "Technical Proposal for the Phase-II Upgrade of the CMS Detector", CMS Technical Design Report CERN-LHCC-2015-010, CMS-TDR-15-02, 2013.

[118] A. F. Kovarik, "On the automatic registration of α-particles, β-particles and γ-ray and X-ray pulses", *Phys. Rev.* **13** (1919) 272.

[119] CMS Collaboration, "The CMS trigger system", *JINST* **12** (2017) P01020, `doi:10.1088/1748-0221/12/01/P01020`, `arXiv:1609.02366`.

[120] C. Collaboration, "The Phase-2 Upgrade of the CMS L1 Trigger Interim Technical Design Report", Technical Report CERN-LHCC-2017-013. CMS-TDR-017, Geneva, 2017.

[121] CMS Collaboration, "Level-1 Muon Trigger performance in 2017 data and comparison with the legacy muon trigger system", CMS Detector Performance Summary CMS-DP-2017-041, 2017.

[122] CMS Collaboration, "The CMS Barrel Muon Trigger upgrade", *JINST* (2017) C01095, `doi:10.1088/1748-0221/12/01/C01095`. *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP 2016): Karlsruhe, Germany, September 26–30, 2016.*

[123] D. Acosta et al., "The CMS Modular Track Finder boards, MTF6 and MTF7", *JINST* **8** (2013) C12034, `doi:10.1088/1748-0221/8/12/C12034`. *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP13): Perugia, Italy, September 23–27, 2013.*

[124] W. M. Zabolotny and A. Byszuk, "Algorithm and implementation of muon trigger and data transmission system for barrel-endcap overlap region of the CMS detector", *JINST* **11** (2016) C03004, `doi:10.1088/1748-0221/11/03/C03004`. *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP15): Lisbon, Portugal, September 28–October 2, 2015.*

[125] CMS Collaboration, "Triggering on electrons, jets and tau leptons with the CMS upgraded Calorimeter Trigger for the LHC Run II", *JINST* **11** (2016), no. 02, C02008, `doi:10.1088/1748-0221/11/02/C02008`. *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP15): Lisbon, Portugal, September 28–October 2, 2015.*

[126] CMS Collaboration, "Tau lepton trigger and identification at CMS in Run-2", *Nucl. Part. Phys. Proc.* **287-288** (2017) 107, `doi:10.1016/j.nuclphysbps.2017.03.055`. *Proceedings, 14th International Workshop on Tau Lepton Physics (TAU 2016): Beijing, China, September 19–23, 2016.*

[127] CMS Collaboration, "Level-1 jets and sums trigger performance", CMS Detector Performance Summary CMS-DP-2016-034, 2016.

[128] K. Compton et al., "The MP7 and CTP-6: Multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS", *JINST* **7** (2012) C12024, `doi:10.1088/1748-0221/7/12/C12024`. *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP12): Oxford, UK, September 17–21, 2012.*

[129] CMS Collaboration, "The CMS Calorimeter Trigger for LHC Run II", *J. Phys. Conf. Ser.* **928** (2017), no. 1, 012004, `doi:10.1088/1742-6596/928/1/012004`. *Proceedings, 17th International Conference on Calorimetry in Particle Physics (CALOR 2016): Daegu, Republic of Korea, May 15–20, 2016.*

[130] A. Zabi et al., "The CMS Level-1 Calorimeter Trigger for the LHC Run II", *JINST* **12** (2017), no. 01, C01065, `doi:10.1088/1748-0221/12/01/C01065`. *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP 2016): Karlsruhe, Germany, September 26–30, 2016.*

[131] CMS Collaboration, C. Amendola, "The CMS Level-1 tau lepton and vector boson fusion triggers for the LHC Run II", in *Proceedings, 2017 European Physical Society Conference on High Energy Physics (EPS-HEP 2017): Venice, Italy, July 5–12, 2017*, volume EPS-HEP2017, p. 773. 2017. `doi:10.22323/1.314.0773`.

[132] CMS Collaboration, "Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV", *JINST* **7** (2012) P10002, `doi:10.1088/1748-0221/7/10/P10002`, `arXiv:1206.4071`.

[133] CMS Collaboration, "HLT Plots Approved for CHEP 2016, San Fransisco". Online at `https://twiki.cern.ch/twiki/bin/view/CMSPublic/HLTplotsCHEP2016`, 2008. Accessed 23 July 2019.

[134] A. Buckley et al., "General-purpose event generators for LHC physics", *Phys. Rept.* **504** (2011) 145, `doi:10.1016/j.physrep.2011.03.005`, `arXiv:1101.2599`.

[135] M. Bahr et al., "Herwig++ physics and manual", *Eur. Phys. J. C* **58** (2008) 639, `doi:10.1140/epjc/s10052-008-0798-9`, `arXiv:0803.0883`.

[136] T. Sjöstrand, S. Ask, J. R. Christiansen et al., "An introduction to PYTHIA 8.2", *Comput. Phys. Commun.* **191** (2015) 159, `doi:10.1016/j.cpc.2015.01.024`, `arXiv:1410.3012`.

[137] J. Alwall et al., "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations", *JHEP* **07** (2014) 079, `doi:10.1007/JHEP07(2014)079`, `arXiv:1405.0301`.

[138] P. Nason, "A New method for combining NLO QCD with shower Monte Carlo algorithms", *JHEP* **11** (2004) 040, `doi:10.1088/1126-6708/2004/11/040`, `arXiv:hep-ph/0409146`.

[139] S. Frixione et al., "Matching NLO QCD computations with Parton Shower simulations: the POWHEG method", *JHEP* **11** (2007) 070, `doi:10.1088/1126-6708/2007/11/070`, `arXiv:0709.2092`.

[140] S. Alioli et al., "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX", *JHEP* **06** (2010) 043, `doi:10.1007/JHEP06(2010)043`, `arXiv:1002.2581`.

[141] J. Alwall et al., "A standard format for Les Houches event files", in *Monte Carlos for the LHC: A Workshop on the Tools for LHC Event Simulation (MC4LHC), Geneva, Switzerland, July 17–16, 2006*, volume 176, p. 300. 2007. `arXiv:hep-ph/0609017`. `doi:10.1016/j.cpc.2006.11.010`.

[142] S. Höche, "Introduction to parton-shower event generators", in *Proceedings, Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders (TASI 2014): Boulder, Colorado, June 2–27, 2014*, p. 235. 2015. `arXiv:1411.4085`. `doi:10.1142/9789814678766_0005`.

[143] C. Degrande et al., "UFO – The Universal FeynRules Output", *Comput. Phys. Commun.* **183** (2012) 1201, `doi:10.1016/j.cpc.2012.01.022`, `arXiv:1108.2040`.

[144] V. N. Gribov and L. N. Lipatov, "Deep inelastic e p scattering in perturbation theory", *Sov. J. Nucl. Phys.* **15** (1972) 438.

[145] L. N. Lipatov, "The parton model and perturbation theory", *Sov. J. Nucl. Phys.* **20** (1975) 94.

[146] Y. L. Dokshitzer, "Calculation of the structure functions for deep inelastic scattering and e$^+$ e$^-$ annihilation by perturbation theory in quantum chromodynamics.", *Sov. Phys. JETP* **46** (1977) 641.

[147] G. Altarelli and G. Parisi, "Asymptotic freedom in parton language", *Nucl. Phys. B* **126** (1977) 298–318, `doi:10.1016/0550-3213(77)90384-4`.

[148] NNPDF Collaboration, "Parton distributions for the LHC Run II", *JHEP* **04** (2015) 040, `doi:10.1007/JHEP04(2015)040`, `arXiv:1410.8849`.

[149] Particle Data Group Collaboration, "Review of particle physics", *Chin. Phys. C* **40** (2016), no. 10, 100001, `doi:10.1088/1674-1137/40/10/100001`.

[150] T. D. Gottschalk, "Backwards evolved initial state parton showers", *Nucl. Phys. B* **277** (1986) 700–738, `doi:10.1016/0550-3213(86)90465-7`.

[151] B. R. Webber.

[152] GEANT4 Collaboration, "GEANT4: A simulation toolkit", *Nucl. Instrum. Meth. A* **506** (2003) 250, `doi:10.1016/S0168-9002(03)01368-8`.

[153] J. Allison et al., "GEANT4 developments and applications", *IEEE Trans. Nucl. Sci.* **53** (2006) 270, `doi:10.1109/TNS.2006.869826`.

[154] CMS Collaboration, S. Sekmen, "Recent developments in CMS fast simulation", in *Proceedings, 38th International Conference on High Energy Physics (ICHEP 2016): Chicago, IL, USA, August 3–10, 2016*, volume ICHEP2016, p. 181. 2016. `arXiv:1701.03850. doi:10.22323/1.282.0181`.

[155] CMS Collaboration, "Particle-flow reconstruction and event description with the CMS detector", *JINST* **12** (2017) P10003, `doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965`.

[156] CMS Collaboration, "Mini-AOD: A new analysis data format for CMS", *J. Phys. Conf. Ser.* **664** (2015) 072052, `doi:10.1088/1742-6596/664/7/072052, arXiv:1702.04685`. *Proceedings, 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13–17, 2015.*

[157] ALEPH Collaboration, "Performance of the ALEPH detector at LEP", *Nucl. Instrum. Meth. A* **360** (1995) 481, `doi:10.1016/0168-9002(95)00138-7`.

[158] ATLAS Collaboration, "Jet reconstruction and performance using particle flow with the ATLAS Detector", *Eur. Phys. J. C* **77** (2017), no. 7, 466, `doi:10.1140/epjc/s10052-017-5031-2, arXiv:1703.10485`.

[159] CMS Collaboration, "Description and performance of track and primary-vertex reconstruction with the CMS tracker", *JINST* **9** (2014), no. 10, P10009, `doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569`.

[160] K. Rose, "Deterministic annealing for clustering, compression, classification, regression, and related optimization problems", in *Proceedings, 1998 International Conference on Electronics, Circuits and Systems (IEEE 1998): Lisboa, Portugal, September 7–10, 1998*, p. 2210. 1998. `doi:10.1109/5.726788`.

[161] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_T jet clustering algorithm", *JHEP* **04** (2008) 063, `doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189`.

[162] M. Cacciari, G. P. Salam, and G. Soyez, "FastJet user manual", *Eur. Phys. J. C* **72** (2012) 1896, `doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097`.

[163] CMS Collaboration, "Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV", *JINST* **13** (2018) P06015, `doi:10.1088/1748-0221/13/06/P06015`, `arXiv:1804.04528`.

[164] R. Fruhwirth, "Application of Kalman filtering to track and vertex fitting", *Nucl. Instrum. Meth. A* **262** (1987) 444, `doi:10.1016/0168-9002(87)90887-4`.

[165] CMS Collaboration, "Performance of missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV using the CMS detector", CMS Physics Analysis Summary CMS-PAS-JME-17-001, 2018.

[166] CMS Collaboration, "Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV", *JINST* **10** (2015) P06005, `doi:10.1088/1748-0221/10/06/P06005`, `arXiv:1502.02701`.

[167] CMS Collaboration, "Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV", *JINST* **12** (2017) P02014, `doi:10.1088/1748-0221/12/02/P02014`, `arXiv:1607.03663`.

[168] D. Bertolini et al., "Pileup per particle identification", *JHEP* **10** (2014) 059, `doi:10.1007/JHEP10(2014)059`, `arXiv:1407.6013`.

[169] CMS Collaboration, M. Stoever, "Jet energy calibrations at CMS experiment with 13 TeV collisions", in *Proceedings, 2017 European Physical Society Conference on High Energy Physics (EPS-HEP 2017): Venice, Italy, July 5–12, 2017*, p. 805. 2017. `doi:10.22323/1.314.0805`.

[170] CMS Collaboration, "Jet algorithms performance in 13 TeV data", CMS Physics Analysis Summary CMS-PAS-JME-16-003, 2017.

[171] CMS Collaboration, "Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV", *JINST* **13** (2018) P05011, `doi:10.1088/1748-0221/13/05/P05011`, `arXiv:1712.07158`.

[172] M. Cacciari and G. P. Salam, "Pileup subtraction using jet areas", *Phys. Lett. B* **659** (2008) 119, `doi:10.1016/j.physletb.2007.09.077`, `arXiv:0707.1378`.

[173] CMS Collaboration, "Identification of b-quark jets with the CMS experiment", *JINST* **8** (2013) P04013, `doi:10.1088/1748-0221/8/04/P04013`, `arXiv:1211.4462`.

[174] CMS Collaboration, "Reconstruction and identification of τ lepton decays to hadrons and ν_τ at CMS", *JINST* **11** (2016) P01019, doi:10.1088/1748-0221/11/01/P01019, arXiv:1510.07488.

[175] CMS Collaboration, "Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV", *JINST* **13** (2018) P10005, doi:10.1088/1748-0221/13/10/P10005, arXiv:1809.02816.

[176] R. Fruhwirth, W. Waltenberger, and P. Vanlaer, "Adaptive vertex fitting", *J. Phys. G* **34** (2007) N343, doi:10.1088/0954-3899/34/12/N01.

[177] ATLAS and CMS Collaborations, The LHC Higgs Combination Group, "Procedure for the LHC Higgs boson search combination in Summer 2011", Technical Report CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-11, 2011.

[178] T. Junk, "Confidence level computation for combining searches with small statistics", *Nucl. Instrum. Meth. A* **434** (1999) 435, doi:10.1016/S0168-9002(99)00498-2, arXiv:hep-ex/9902006.

[179] A. L. Read, "Presentation of search results: The CL$_s$ technique", *J. Phys. G* **28** (2002) 2693, doi:10.1088/0954-3899/28/10/313.

[180] J. Sonneveld, "Limit setting and statistics". Presentation at *The first MadAnalysis 5 workshop on LHC recasting @ Korea*: Seoul, Korea, August 20–27, 2017. Online at https://indico.cern.ch/event/637941/timetable/#7-limit-setting-and-statistics, 2017. Accessed 5 October 2019.

[181] G. Cowan et al., "Asymptotic formulae for likelihood-based tests of new physics", *Eur. Phys. J. C* **71** (2011) 1554, doi:10.1140/epjc/s10052-011-1554-0, arXiv:1007.1727. [Erratum: *Eur. Phys. J.C* **73** (2013) 2501].

[182] A. Read, "Linear interpolation of histograms", *Nucl. Instrum. Meth. Res. A* **425** (1999) 357.

[183] R. J. Barlow and C. Beeston, "Fitting using finite Monte Carlo samples", *Comput. Phys. Commun.* **77** (1993) 219, doi:10.1016/0010-4655(93)90005-W.

[184] P. Artoisenet et al., "Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations", *JHEP* **03** (2013) 015, doi:10.1007/JHEP03(2013)015, arXiv:1212.3460.

[185] T. Ježo et al., "An NLO+PS generator for $t\bar{t}$ and Wt production and decay including non-resonant and interference effects", *Eur. Phys. J. C* **76** (2016) 691, `doi:10.1140/epjc/s10052-016-4538-2`, `arXiv:1607.04538`.

[186] S. Frixione, P. Nason, and G. Ridolfi, "A Positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction", *JHEP* **09** (2007) 126, `doi:10.1088/1126-6708/2007/09/126`, `arXiv:0707.3088`.

[187] R. Frederix and S. Frixione, "Merging meets matching in MC@NLO", *JHEP* **12** (2012) 061, `doi:10.1007/JHEP12(2012)061`, `arXiv:1209.6215`.

[188] S. Alioli et al., "NLO single-top production matched with shower in POWHEG: s- and t-channel contributions", *JHEP* **09** (2009) 111, `doi:10.1007/JHEP02(2010)011`, `arXiv:0907.4076`. [Erratum: *JHEP* **10** (2009) 011].

[189] E. Re, "Single-top Wt-channel production matched with parton showers using the POWHEG method", *Eur. Phys. J. C* **71** (2011) 1547, `doi:10.1140/epjc/s10052-011-1547-z`, `arXiv:1009.2450`.

[190] J. Alwall et al., "Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions", *Eur. Phys. J. C* **53** (2008) 473, `doi:10.1140/epjc/s10052-007-0490-5`, `arXiv:0706.2569`.

[191] CMS Collaboration, "Event generator tunes obtained from underlying event and multiparton scattering measurements", *Eur. Phys. J. C* **76** (2016) 155, `doi:10.1140/epjc/s10052-016-3988-x`, `arXiv:1512.00815`.

[192] CMS Collaboration, "Investigations of the impact of the parton shower tuning in PYTHIA 8 in the modelling of $t\bar{t}$ at $\sqrt{s} = 8$ and 13 TeV", CMS Physics Analysis Summary CMS-PAS-TOP-16-021, 2016.

[193] NNPDF Collaboration, "Unbiased global determination of parton distributions and their uncertainties at NNLO and at LO", *Nucl. Phys. B* **855** (2012) 153, `doi:10.1016/j.nuclphysb.2011.09.024`, `arXiv:1107.2652`.

[194] M. Lotti, "Trigger efficiency measurement in the charged Higgs boson $H^+ \rightarrow \tau\nu$ analysis", Master's thesis, University of Helsinki, Faculty of Science, Department of Physics, 2017.

[195] M. Kortelainen, "Search for a light charged Higgs boson in the CMS experiment in pp collisions at $\sqrt{(s)} = 7$ TeV". PhD thesis, University of Helsinki, Faculty of Science, Department of Physics, 2012.

[196] E. Pekkarinen, "Data-driven measurement of the background with misidentified tau leptons in a search for charged Higgs bosons", Master's thesis, Aalto University, School of Science, 2015.

[197] M. Czakon and A. Mitov, "Top++: A program for the calculation of the top-pair cross-section at hadron colliders", *Comput. Phys. Commun.* **185** (2014) 2930, `doi:10.1016/j.cpc.2014.06.021`, `arXiv:1112.5675`.

[198] N. Kidonakis, "Top Quark Production", in *Proceedings, Helmholtz International Summer School on Physics of Heavy Quarks and Hadrons (HQ 2013), Dubna, Russia*, p. 139. 2014. `arXiv:1311.0283`. `doi:10.3204/DESY-PROC-2013-03/Kidonakis`.

[199] P. Kant et al., "HatHor for single top-quark production: Updated predictions and uncertainty estimates for single top-quark production in hadronic collisions", *Comput. Phys. Commun.* **191** (2015) 74, `doi:10.1016/j.cpc.2015.02.001`, `arXiv:1406.4403`.

[200] K. Melnikov and F. Petriello, "Electroweak gauge boson production at hadron colliders through $O(\alpha_s^2)$", *Phys. Rev. D* **74** (2006) 114017, `doi:10.1103/PhysRevD.74.114017`, `arXiv:hep-ph/0609070`.

[201] J. M. Campbell, R. K. Ellis, and C. Williams, "Vector boson pair production at the LHC", *JHEP* **07** (2011) 018, `doi:10.1007/JHEP07(2011)018`, `arXiv:1105.0020`.

[202] CMS Collaboration, "Measurement of the $t\bar{t}$ production cross section, the top quark mass, and the strong coupling constant using dilepton events in pp collisions at $\sqrt{s} = 13$ TeV", *Eur. Phys. J. C* **79** (2019) 368, `doi:10.1140/epjc/s10052-019-6863-8`, `arXiv:1812.10505`.

[203] CMS Collaboration, "Measurements of $t\bar{t}$ differential cross sections in proton-proton collisions at $\sqrt{s} = 13$ TeV using events containing two leptons", *JHEP* **02** (2019) 149, `doi:10.1007/JHEP02(2019)149`, `arXiv:1811.06625`.

[204] CMS Collaboration, "Search for neutral MSSM Higgs bosons decaying to a pair of tau leptons in pp collisions", *JHEP* **10** (2014) 160, `doi:10.1007/JHEP10(2014)160`, `arXiv:1408.3316`.

[205] CMS Collaboration, "An embedding technique to determine $\tau\tau$ backgrounds in proton-proton collision data", *JINST* **14** (2019), no. 06, P06032, `doi:10.1088/1748-0221/14/06/P06032`, `arXiv:1903.01216`.

[206] CMS Collaboration, "Measurement of Higgs boson production and decay to the $\tau\tau$ final state", CMS Physics Analysis Summary CMS-PAS-HIG-18-032, 2019.

[207] J. Butterworth et al., "PDF4LHC recommendations for LHC Run II", *J. Phys. G* **43** (2016) 023001, `doi:10.1088/0954-3899/43/2/023001`, `arXiv:1510.03865`.

[208] ATLAS Collaboration, "Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 13$ TeV with the ATLAS detector at the LHC", *Phys. Rev. Lett.* **117** (2016) 182002, `doi:10.1103/PhysRevLett.117.182002`, `arXiv:1606.02625`.

[209] CMS Collaboration, "CMS Luminosity Measurements for the 2016 Data Taking Period", CMS Physics Analysis Summary CMS-PAS-LUM-17-001, 2017.

[210] M. Flechl et al., "Improved cross-section predictions for heavy charged Higgs boson production at the LHC", *Phys. Rev. D* **91** (2015) 075015, `doi:10.1103/PhysRevD.91.075015`, `arXiv:1409.5615`.

[211] C. Degrande et al., "Heavy charged Higgs boson production at the LHC", *JHEP* **10** (2015) 145, `doi:10.1007/JHEP10(2015)145`, `arXiv:1507.02549`.

[212] S. Dittmaier et al., "Charged-Higgs-boson production at the LHC: NLO supersymmetric QCD corrections", *Phys. Rev. D* **83** (2011) 055005, `doi:10.1103/PhysRevD.83.055005`, `arXiv:0906.2648`.

[213] E. L. Berger et al., "Associated production of a top quark and a charged Higgs boson", *Phys. Rev. D* **71** (2005) 115012, `doi:10.1103/PhysRevD.71.115012`, `arXiv:hep-ph/0312286`.

by Keti Mirziashvili