



Faculté
des
Sciences



Search for promptly decaying heavy neutral leptons in proton-proton collision events with trilepton final states using the CMS detector

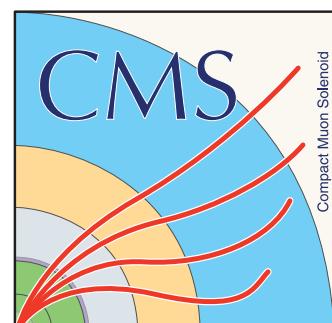
Thesis presented by Liam WEZENBEEK

in fulfilment of the requirements of the PhD Degree in Sciences (ULB - “Doctorat en Sciences”) and in Sciences: Physics (UGent)
Academic year 2023-2024

Supervisors: Professor Barbara CLERBAUX (Université libre de Bruxelles)
and Professor Didar DOBUR (Universiteit Gent)

Thesis jury :

Philippe SMET (UGent, Chair)
Laurent THOMAS (ULB, Secretary)
Archisman GHOSH (UGent)
Richard RUIZ (IFJ PAN)
Livia SOFFI (INFN Rome)



Abstract

This dissertation presents two analyses that search for signs of physics beyond the standard model, which describes all known particles and their interactions at the most fundamental level. While experimental evidence has proven the theory to be highly robust, it has also shown that the model is incomplete. A number of open questions remain unsolved within the standard model framework. In order to address these open questions, new models that serve as extensions to the standard model have been proposed.

Both searches described in this work use a dataset recorded by the CMS experiment between 2016 and 2018, corresponding to a total integrated luminosity of 137 fb^{-1} . The first probes this dataset for signs of electroweak superpartner particles that would point to the existence of a supersymmetry. This is done using a mixture of classic cut-based search regions and novel neural network architectures. A wide range of simplified models that propose a final state of two same-sign leptons or three leptons or more are considered in the design and interpretation but this dissertation will focus on the contributions of the author to the final states with a hadronically decayed tau lepton. No excess of events over the predicted background is observed. Stringent limits on the cross sections of all considered models are placed.

The second analysis looks for hints of promptly decaying Dirac and Majorana heavy neutral leptons (HNL) in events with three leptons of which at most one is allowed to be a τ lepton. The existence of such particles could explain why the standard model neutrinos have such low masses or why there is more matter than anti-matter in the universe. Sterile neutrino masses in the range of 10 GeV to 1.5 TeV are considered in models that pose the addition of a single HNL to the standard model. In these models, the HNL couples exclusively to a single standard model neutrino flavour. The search scrutinizes the same 2016-2018 dataset using a combination of cut-based search regions and boosted decision trees. No excess of events over the predicted background is observed. Improved limits on the interaction strength as a function of HNL mass over the state-of-the-art exclusion limits are produced in scenarios with exclusive coupling to one of the light flavour standard model neutrinos. The analysis also provides the first such limit for scenarios with exclusive tau neutrino coupling from the LHC.

0. Abstract

Samenvatting

Deze thesis presenteert twee analyses die zoeken naar tekenen van fysica buiten het standaardmodel, dat alle gekende deeltjes en hun interacties op het meest fundamentele niveau beschrijft. Hoewel experimenteel bewijs heeft aangetoond dat de theorie enorm robust is, heeft het ook aangetoond dat het model onvolledig is. Binnen het kader van het standaardmodel blijven er een aantal open vragen onopgelost. Om deze open vragen in te vullen, werden nieuwe modellen voorgesteld die het standaardmodel verder uitbreiden.

Beide analyses die in dit werk worden beschreven, maken gebruik van een dataset die is opgenomen door het CMS-experiment tussen 2016 en 2018, wat overeenkomt met een totale luminositeit van 137 fb^{-1} . De eerste analyse gaat in deze dataset op zoek naar tekenen van nieuwe elektrozwakke superpartner deeltjes die wijzen op het bestaan van een supersymmetrie met behulp van zowel cut-gebaseerde zoekregio's als nieuwe neurale netwerkarchitecturen. Een breed scala aan vereenvoudigde modellen die eindigen in het verval naar een toestand van meerdere leptonen wordt onderzocht in de analyse. Desondanks zal deze thesis zich focussen op slechts één model dat zal dienen als ondersteunend materiaal om de bijdragen van de auteur in eindtoestanden met een hadronisch vervallen tau lepton te beschrijven. De resultaten tonen een goede overeenkomst tussen waargenomen data en voorspelde achtergrond. Strenge limieten worden geplaatst op de crosssecties van de beschouwde modellen.

De tweede analyse zoekt naar de aanwezigheid van Dirac- en Majorana-type zware neutrale leptonen in evenementen met drie leptonen, waarvan er maximaal één een τ lepton mag zijn. Het bestaan van zulke neutrale leptonen zou kunnen verklaren waarom de neutrino massa's in het standaardmodel zo laag zijn of waarom er meer materie dan antimaterie in het universum is. Steriele neutrino's met massa's in het bereik van 10 GeV en 1.5 TeV worden in rekening genomen in modellen die slechts één extra zulk neutrino in het standaardmodel voorspellen en waarbij dit deeltje enkel mag koppelen met een van de drie standaardmodel neutrino's per keer. De zoektocht onderzoekt dezelfde dataset van 2016-2018 met behulp van een combinatie van cut-gebaseerde zoekregio's en boosted decision trees. Ook hier komt de data overeen met de voorspelde achtergrond. In vergelijking met voorgaande analyses worden er verbeterde limieten op de koppelingssterkte als functie van steriele neutrino-massa geplaatst in koppelingsscenario's waar het steriele neutrino koppelt met een electron- of muonneutrino. Deze analyse levert ook de eerste limieten voor scenarios met exclusieve koppeling tussen τ -neutrino en steriel neutrino aan de LHC.

0. Samenvatting

Résumé

Cette thèse présente deux analyses qui recherchent des signes de nouvelle physique au-delà du modèle standard. Ce dernier décrit toutes les particules connues et leurs interactions au niveau le plus fondamental. Alors que les mesures expérimentales ont prouvé que la théorie était très robuste, elles ont également montré que le modèle est incomplet. Un certain nombre de questions ouvertes restent non résolues dans le cadre du modèle standard. Afin de répondre à ces questions ouvertes, de nouveaux modèles qui servent d'extensions au modèle standard ont été proposés.

Les deux recherches décrites dans ce travail utilisent un jeu de données enregistré par l'expérience CMS entre 2016 et 2018, correspondant à une luminosité totale de 137 fb^{-1} . Le premier sonde cet ensemble de données à la recherche de signes de nouvelles particules superpartenaires électrofaibles qui indiquerait l'existence d'une supersymétrie. Cela se fait en utilisant un mélange de régions de recherche classiques basées sur des coupures et de nouvelles architectures de réseaux neuronaux. Une large gamme de modèles simplifiés qui proposent un état final de deux leptons de même signe ou de trois leptons ou plus sont pris en compte dans la conception et l'interprétation de l'analyse, mais cette thèse se concentrera sur les contributions de l'auteur aux états finaux avec un lepton tau se désintégrant hadroniquement. Aucun excès d'événements par rapport au bruit de fond prévu n'est observé. Des limites strictes sur les sections efficaces de tous les modèles considérés sont placées.

La deuxième analyse recherche des indices de désintégration rapide des leptons neutres lourds (HNL) de Dirac et de Majorana dans des événements avec trois leptons dont au plus un est autorisé à être un lepton tau. L'existence de telles particules pourrait expliquer pourquoi les neutrinos du modèle standard ont des masses si faibles ou pourquoi il y a plus de matière que d'antimatière dans l'univers. Les masses de neutrinos stériles comprises entre 10 GeV et 1,5 TeV sont prises en compte dans les modèles qui imposent l'ajout d'un seul HNL qui se couple exclusivement au neutrino de l'électron, au neutrino du muon ou au neutrino du tau du modèle standard. Cette recherche examine le même ensemble de données 2016-2018 en utilisant une combinaison de régions de recherche basées sur des coupures et des arbres de décision améliorés. Aucun excès d'événements par rapport au bruit de fond prévu n'est observé. Des limites améliorées sur la force d'interaction en fonction de la masse de HNL par rapport à l'état de l'art sont produites dans des scénarios avec un couplage exclusif à l'un des neutrinos du modèle standard de saveur légère. L'analyse fournit également la première limite de ce type pour les scénarios avec couplage exclusif des neutrinos du tau du LHC.

0. Résumé

Acknowledgments

As I approach my public defense, I have been taking time to reflect on the road that has led me here. This exhausting, yet incredibly satisfying journey is not one I undertook in solitude, but one shaped by the support of numerous individuals who have contributed in some way or another to the completion of this work. To these people, I would like to express my sincere gratitude below.

I would like to start by thanking my promotores, prof. dr. Didar Dobur and prof. dr. Barbara Clerbaux, for their continued guidance in this project. Didar, thank you for all the opportunities that allowed me to enter the field. The work I have done with you, which started a long time ago with an internship at CERN, have made me into the researcher I am today. Barbara, thank you for maintaining a realistic outlook in the project. Without your guidance, I would have got stuck in many unfeasible timelines and ideas. Additionally, I would like to thank prof. dr. Dirk Ryckbosch. Your engaging lectures have sparked an interest in me that fueled this whole experience.

During the past five years, I was blessed with wonderful colleagues who not only engaged in constructive discussions with me but also brought joy to the workplace. Thank you Martina, Willem, Illia, Tom, Basile, Daniele and Marek for taking me in as the new kid and raising me with the needed knowledge to perform my own analysis. Thank you Joscha, for being the steadfast pillar of expertise that made my life so much easier in these last few years. Thank you to the younger generation of colleagues with whom I shared an office – Luka, Niels and David – for continually challenging me on topics I thought I had already mastered and turning our office into a playground filled with ridiculous pastimes. A special thanks goes to Gianny, not just a colleague but a close friend who always had my back. Mental breakdowns were always more fun when there were two of us hiding below the table. I would also like to thank two colleagues outside the CMS group who I have grown very fond of: Stef and Tom. You were always the first people I could go to for much needed breaks or discussions.

Of course there are also many people that have not made direct scientific contributions to this work but whose support was essential to its completion. Firstly, I would like to thank my oldest band of friends: Alex, Injas, Janne, Remy, Wouter, Jasper, Arne and Peter. Thank you for all the good times and for getting me out of the house when I was about to get lost in work again. Thank you, Lieven, for being one of the most loyal friends I have. I could always count on you in times of need. I would also like to express my gratitude to the people I have met through the PhD community: Kaj, Elien, Seppe, Tiago, Prudence and many more. Although

0. Résumé

we only met last year, you have all grown dear to me. Hearing about your PhD experiences has always helped to put my own issues into a broader context.

Next, I would like to thank two people that are very important to me – my parents. Your contributions to this work can not be emphasized enough. You have always given me the freedom and resources to become the person I am today and you are always ready to help me when I need it. And last but not least, thank you, Margaux, my better half. Although you have not been in my life for that long, you have made the largest impact. Thank you for your unconditional support and for keeping this sinking ship afloat.

Liam Wezenbeek

Ghent, 2023

Contents

Abstract

Samenvatting

Résumé

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | The standard model | 5 |
| 2.1 | Symmetry in nature | 5 |
| 2.2 | A phenomenological description of the standard model | 6 |
| 2.2.1 | Fundamental interactions and gauge bosons | 7 |
| 2.2.2 | Fermions | 9 |
| 2.3 | Some concepts from Quantum Field Theory | 11 |
| 2.3.1 | The standard model | 12 |
| 2.3.2 | Spontaneous symmetry breaking | 13 |
| 2.3.3 | Feynman diagrams | 15 |
| 2.4 | Limitations of the standard model | 15 |
| 2.4.1 | The Hierarchy Problem | 15 |
| 2.4.2 | Neutrino mass | 16 |
| 2.4.3 | Dark matter | 16 |
| 2.4.4 | Matter-antimatter asymmetry | 18 |
| 3 | Extending the standard model | 19 |
| 3.1 | Supersymmetry | 19 |
| 3.1.1 | Solving the hierarchy problem | 19 |
| 3.1.2 | Supersymmetry and soft breaking | 20 |
| 3.1.3 | Minimal Supersymmetric Standard Model | 21 |
| 3.1.4 | R-parity | 22 |
| 3.1.5 | Supersymmetry at the LHC | 23 |
| 3.2 | Neutrino Physics | 25 |
| 3.2.1 | Introduction | 25 |
| 3.2.2 | Neutrino Oscillations | 26 |
| 3.2.3 | Dirac, Majorana and Weyl fermions | 27 |
| 3.2.4 | Massive neutrinos | 28 |
| 3.2.5 | The Seesaw Mechanism | 29 |
| 3.2.6 | Low-scale models | 31 |
| 3.2.7 | State-of-the-art constraints | 31 |

| | |
|--|-----------|
| 4 The CMS experiment | 37 |
| 4.1 The Large Hadron Collider | 37 |
| 4.2 The Compact Muon Solenoid | 40 |
| 4.2.1 The CMS coordinate system | 40 |
| 4.2.2 The CMS detector | 42 |
| 4.2.3 Trigger system | 47 |
| 4.3 Object reconstruction in CMS | 48 |
| 4.3.1 Tracking and vertex reconstruction | 48 |
| 4.3.2 Calorimeter clustering | 49 |
| 4.3.3 Particle Flow | 49 |
| 4.3.4 Muon reconstruction | 49 |
| 4.3.5 Electron reconstruction | 50 |
| 4.3.6 Jet clustering | 51 |
| 4.3.7 Heavy Flavor identification | 52 |
| 4.3.8 Tau Leptons | 52 |
| 4.3.9 Missing energy | 55 |
| 4.4 Event simulation | 55 |
| 4.4.1 Matrix element-level calculations | 56 |
| 4.4.2 Decays and hadronisation | 57 |
| 4.4.3 Detector simulation | 57 |
| 4.5 Statistical tools | 58 |
| 5 Triggers at CMS | 61 |
| 5.1 Personal Foreword | 61 |
| 5.2 Introduction | 62 |
| 5.3 Triggers in Run II | 62 |
| 5.3.1 L1 Trigger | 62 |
| 5.3.2 High Level Trigger | 63 |
| 5.3.3 Anatomy of tau triggers | 65 |
| 5.4 Triggers in Run III | 67 |
| 5.5 Tau Triggers in Run III | 68 |
| 5.5.1 Improved L2 sequence | 68 |
| 5.5.2 DeepTau@HLT | 69 |
| 5.5.3 Final strategy and performance | 72 |
| 5.6 Conclusions | 74 |
| 6 A search for supersymmetry in a multilepton final state | 77 |
| 6.1 Personal foreword | 77 |
| 6.2 Introduction | 78 |
| 6.3 Signal Models | 79 |
| 6.4 Data and event simulation | 80 |
| 6.4.1 Collision data | 80 |
| 6.4.2 Event simulation | 80 |
| 6.5 Object selection and identification | 81 |
| 6.5.1 Muons | 81 |
| 6.5.2 Electrons | 83 |
| 6.5.3 Tau Leptons | 84 |

| | | |
|----------|--|------------|
| 6.5.4 | Jets and b-tagging | 87 |
| 6.5.5 | Missing transverse energy | 87 |
| 6.6 | Analysis strategy and event selection | 88 |
| 6.7 | Background Estimation | 92 |
| 6.7.1 | Nonprompt background processes | 93 |
| 6.7.2 | Prompt backgrounds | 98 |
| 6.7.3 | Photon conversions | 104 |
| 6.8 | Systematic uncertainties | 104 |
| 6.9 | Results | 107 |
| 6.10 | Interpretation | 111 |
| 6.11 | Summary | 114 |
| 7 | A search for heavy neutral leptons in a multilepton final state | 117 |
| 7.1 | Personal foreword | 117 |
| 7.2 | Introduction | 117 |
| 7.3 | Signal model | 118 |
| 7.3.1 | Simulated samples | 120 |
| 7.3.2 | Signal compression | 123 |
| 7.4 | Data and background simulation | 124 |
| 7.4.1 | Collision data | 124 |
| 7.4.2 | Standard model process simulation | 125 |
| 7.5 | Object Selection | 125 |
| 7.5.1 | Light Leptons | 125 |
| 7.5.2 | Tau Leptons | 127 |
| 7.5.3 | Jets | 129 |
| 7.5.4 | Missing transverse energy | 130 |
| 7.6 | Trigger strategy | 130 |
| 7.6.1 | Strategy | 130 |
| 7.6.2 | Offline thresholds | 131 |
| 7.6.3 | Trigger efficiency | 133 |
| 7.7 | Analysis Strategy | 136 |
| 7.7.1 | Event selection | 138 |
| 7.7.2 | Cut-based search bins | 139 |
| 7.7.3 | Boosted decision trees | 145 |
| 7.7.4 | Combined strategy | 151 |
| 7.8 | Background Estimation | 153 |
| 7.8.1 | Nonprompt background | 153 |
| 7.8.2 | WZ background | 160 |
| 7.8.3 | ZZ background | 165 |
| 7.8.4 | Background from photon conversions | 166 |
| 7.8.5 | Charge misidentification background | 169 |
| 7.9 | Systematic uncertainties | 169 |
| 7.10 | Results | 172 |
| 7.10.1 | Interpretation | 175 |
| 7.10.2 | Exclusive light lepton neutrino coupling | 187 |
| 7.10.3 | Exclusive tau neutrino coupling | 188 |
| 7.11 | Summary and outlook | 191 |

Contents

| | |
|-----------------------|------------|
| 8 Conclusions | 195 |
| 8.1 Summary | 195 |
| 8.2 Outlook | 196 |
| Bibliography | 199 |

List Of Abbreviations

| | |
|--------------|--|
| AMS | Alpha Magnetic Spectrometer |
| ATLAS | A Toroidal LHC Apparatus |
| AUC | area under the curve |
| BBN | Big Bang Nucleosynthesis |
| BDT | boosted decision tree |
| BEH | Brout-Englert-Higgs |
| CERN | European Laboratory for Particle Physics |
| CPU | central processing unit |
| CL | confidence level |
| CMS | Compact Muon Solenoid |
| CMSSW | CMS Software |
| CNN | convolutional neural network |
| CSC | cathode strip chamber |
| CSV | Combined Secondary Vertex |
| DF | different flavour |
| DT | drift tube |
| DPS | detector performance summary |
| DNN | deep neural network |
| DY | Drell–Yan |
| DQM | data quality monitoring |
| ECAL | electromagnetic calorimeter |
| EWSB | electroweak symmetry breaking |

Contents

| | |
|--------------|----------------------------------|
| FCC | Future Circular Collider |
| FO | fakeable object |
| GPU | graphical processing unit |
| GUT | Grand Unified Theory |
| GSF | gaussian-sum filter |
| GMSB | gauge-mediated SUSY breaking |
| HB | HCAL barrel |
| HCAL | hadronic calorimeter |
| HE | HCAL endcap |
| HF | forward HCAL |
| HL | High-Luminosity |
| HLT | high-level trigger |
| HNL | heavy neutral lepton |
| HO | outer HCAL |
| HPS | Hadrons-Plus-Strips |
| ISR | initial state radiation |
| JEC | jet energy corrections |
| JER | jet energy resolution |
| KF | Kalman Filter |
| L1 | level-1 |
| LO | leading order |
| LHC | Large Hadron Collider |
| LINAC | linear accelerator |
| LEP | Large Electron-Positron Collider |
| LSP | lightest supersymmetric particle |
| ME | matrix element |
| MET | missing transverse momentum |
| MC | Monte Carlo |

Contents

| | |
|-------------|--|
| MSSM | Minimal Supersymmetric Standard Model |
| MVA | multivariate analysis |
| NLSP | next-to-lightest supersymmetric particle |
| NLO | next-to-leading order |
| NNLO | next-to-next-to-leading order |
| OS | opposite sign |
| OSSF | opposite sign and same flavour |
| OSDF | opposite sign and different flavour |
| PDF | parton distribution function |
| PF | Particle Flow |
| PFN | Particle Flow Network |
| PS | Proton Synchrotron |
| POG | particle object group |
| QCD | quantum chromodynamics |
| QED | quantum electrodynamics |
| QFT | quantum field theory |
| RF | radiofrequency |
| RMS | root mean squared |
| ROC | receiver operating characteristic |
| RPC | resistive plate chamber |
| SIP | significant impact parameter |
| SLAC | Stanford Linear Accelerator Center |
| SM | standard model of particle physics |
| SPS | Super Proton Synchrotron |
| SSDF | same sign and different flavour |
| SUSY | supersymmetry |
| TIB | tracker inner barrel |
| TOB | tracker outer barrel |

Contents

TID tracker inner disc

TSG trigger studies group

UL ultra-legacy

VBF vector boson fusion

1

Introduction

The observable universe around us consists of vast amounts of matter clumped into fascinating structures. The idea that this matter is compiled from a limited number of fundamental building blocks goes back as far as the 6th century BC in Indian and ancient Greek philosophy [1,2]. In the 20th century, the scientific community made this hypothesis concrete with the development of the standard model of particle physics (SM), which describes all elementary particles and their interactions in a mathematical framework. Today, it has evolved into a highly robust model that superbly describes a wide range of experimental observations and is still being backed up by additional measurements and discoveries. One of the latest such discoveries was the one of the Brout-Englert-Higgs boson by the Compact Muon Solenoid (CMS) and A Toroidal LHC Apparatus (ATLAS) experiments at the European Laboratory for Particle Physics (CERN) in 2012 [3–6]. By observing the – up until then – last undetected particle predicted by the SM, this discovery has further established the robustness of the model.

While the validity of the SM is generally undisputed, there are a few areas that show that it can not be complete. Theorists have devised a plethora of extensions to the SM that can cleverly close these open issues. Experiments like the aforementioned CMS and ATLAS put these theories to the test by performing searches for signs of the new physics models in their detectors. This work contains a description of two such searches to which the author has contributed. The first of these is a search for supersymmetry (SUSY), which adds an additional symmetry to the SM theory along with an accompanying new “superpartner” particle for every SM particle. As the author was not one of the main analysers for this analysis, we will stick to a brief description of parts of that search that are most relevant to the author’s contributions. The search itself looks for signs of electroweak production of superpartners with subsequent decay to fully leptonic final states. In this dissertation, we will limit ourselves to final states containing three charged leptons of which at least one is a hadronically decayed tau.

The second search will be the main focus of this work. It looks for signs of the presence of heavy neutral leptons in proton-proton collision events in the CMS detector. These new neutrinos could explain why the SM neutrinos have such low masses. As such, this search aims to contribute to the neutrino sector by either discovering new heavy neutrinos or by placing exclusion limits on their existence. This is done by studying benchmark models with a single

additional heavy neutrino that is produced in the decay of a W boson. In these benchmark models, there is exclusive coupling of the heavy neutral lepton to one of the SM neutrinos. Final states of three charged leptons are probed in order to search for signs of a heavy neutrino with a mass in the range of 10 GeV to 1.5 TeV. Not only does this search improve on the existing exclusion limits from CMS, it also offers the first results on the exclusive coupling of such new neutrinos to the τ neutrino from any experiment at the Large Hadron Collider (LHC).

To provide context to the analyses and clarify any terms in the last paragraphs that might be unfamiliar to the reader, this dissertation will open with a few words on the theoretical framework. To this end, Chapter 2 will provide an introduction to the SM of particle physics with a discussion on the open questions at the end. Chapter 3 will then go into extension models to the SM that are relevant to the further chapters. All the work performed by the author was done as part of the CMS experiment and uses data, simulation and tools provided by said collaboration. Chapter 4 will therefore cover the architecture of the CMS detector as well as the tools and framework used to reconstruct signatures and perform the analyses. CMS members are expected to participate in service work for the maintenance, upgrades and operation of the detector. Chapter 5 will be dedicated to the contribution of the author to such service work in the context of the CMS trigger system and its preparation for the next years of data-taking. We will finish with the two aforementioned analyses in Chapters 6 and 7. Finally, in Chapter 8 we will take a look back, draw conclusions and use them to look towards the future.

Contributions by the author

The author has made contributions to the tau trigger development and maintenance in CMS, a search for electroweak production of electroweak superpartners and was the main analyser for a search for heavy neutral leptons in multilepton final states. The personal contributions made to these three projects are outlined at the end of Chapters 5, 6 and 7. Below an overview is given of the publications directly contributed to by the author and the presentations given in the scientific community. It should be noted that the paper for the search on heavy neutral leptons is not yet published but is in the final stages of internal review.

Publications

- [7] CMS Collaboration, “Search for electroweak production of charginos and neutralinos in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2022) 147.

In preparation

- CMS Collaboration, “Search for heavy neutral leptons in decays with electrons, muons, and hadronically decaying tau leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”

Talks

- “Search for heavy neutral leptons in multilepton final states in full Run II data”, at CMS EXO Group Workshop 2020, 27–30 Oct 2020
- “Search for electroweak SUSY production in leptonic and hadronic final states with the CMS experiment”, parallel at EPS-HEP 2021 conference, 26–30 Jul 2021
- “Search for EWK SUSY production and HNLs at CMS”, at EOS equinox meeting, 9 Sep 2021

Through the work described in Chapter 5, the author has made contributions to the following detector performance summary (DPS) note:

- [8] CMS Collaboration, “Performance of tau lepton reconstruction at High Level Trigger using 2022 data from the CMS experiment at CERN”, *CDS, CMS-DP-2023-024* (2023).

2

The standard model

All our knowledge on the fundamental building blocks of the universe and their interactions – with the exception of gravity, which currently can not fully be described as a quantum field theory – is contained in the SM of particle physics. It describes the components that make up all the visible matter in the universe, the messenger particles that allow interplay between them and all their properties. When we start stripping this model down to its essence, we see that it is a gauge theory that is described by a select number of symmetry groups [9]. Section 2.1 will first introduce what we mean by gauge theories and will attempt to establish the structure of the SM at its core level.

The full theory and mechanisms behind the SM are extensive and complex, however. To stay within the scope of the work performed by the author, the reader is referred to external references for more details [10–12]. Instead, a phenomenological description of the particle content and gauge interactions that roll out of this gauge theory is provided in Section 2.2. We will go over all the particles contained in the SM and their properties as well as their interplay with the fundamental forces and each other without deriving them from first principle.

Once all the players are on the board, a concise introduction to the theoretical aspects of the SM is provided in Section 2.3 to give some context to later chapters. It is not meant to give a full theoretical foundation of the complete SM but instead to give a better understanding to the reader of topics relevant to the analyses discussed in the last two chapters. It will mainly lead up to discuss the symmetry breaking in the SM that provides particles with their mass.

To finish the chapter, both points of view will be brought together to discuss some of the open issues the SM is currently not able to explain in Section 2.4. Some of these open questions are the driving factor behind the importance of the new physics searches and show the need for experiments like CMS to explore the high-energy frontier in search of answers to them.

2.1 Symmetry in nature

The concept of symmetry, referring to invariance under transformations, is paramount to all laws of nature [13, 14]. Not only does it introduce perceived beauty into our universe, it is

also responsible for a plethora (if not all) of phenomena we observe. Maxwell's equations of electromagnetism follow from symmetry of the electromagnetic field. Einstein's theory of relativity is linked to the symmetry of space-time itself. Emmy Noether showed in 1918 that symmetry and conservation laws are closely linked [15]. For every continuous symmetry, there is a time-independent quantity that is conserved. This quantity is often referred to as a charge. It can be shown that conservation of momentum is a direct result of the translational symmetry of space. Similarly, energy conservation is linked to the translational symmetry of time.

A symmetry is labelled as *global* when its transformations take effect in every point of space. Performing a global rotation during an experiment means that both experiment and observer perform the same rotation. *Local symmetries* or *gauge symmetries* on the other hand are evaluated at a specific point in space. The SM is described by such a symmetry: the invariance of fields under local transformations in certain symmetry groups dictates the laws of fundamental particle physics.

A fundamental symmetry group can be assigned to each of the fundamental forces. While rotational symmetry is easy to imagine, these symmetry groups are somewhat more abstract. The strong force that keeps atomic nuclei together, for example, is described by the $SU(3)$ symmetry group. This is the special unitarity group of degree three, the fundamental representation of which are the 3×3 unitary matrices with determinant one. Similarly, the weak force governing nuclear beta decay is described by the $SU(2)$ group and the electromagnetic force is described by the $U(1)$ group. Together these three form the group $SU(3)_C \times SU(2)_L \times U(1)_Y$ that governs the gauge theory called the SM [10–12].

2.2 A phenomenological description of the standard model

The quantum field theory called the SM describes all fundamental particles as a superposition of quantum fields that are evaluated at every point in space-time. The physical particles that we observe are then resonances of those fields that form real, local particles. All particles in the SM can be split into two main groups depending on the value of an internal degree of freedom of the particle called “*spin*”. It is a quantum number that will determine the behaviour of the particle in a system. Particles with a half-integer value for their spin will follow Fermi-Dirac statistics [16]. This means that the quantum states are distributed over a Fermi-Dirac distribution of energy states where no two particles can occupy the same state. This type of particle is appropriately called a *fermion*. *Bosons* on the other hand are particles with integer values for their spin. They instead follow Bose-Einstein statistics and are not limited to one particle per quantum state. An overview of all fermions and bosons in the SM is given in Figure 2.1. We will go into each of them in the subsections below.

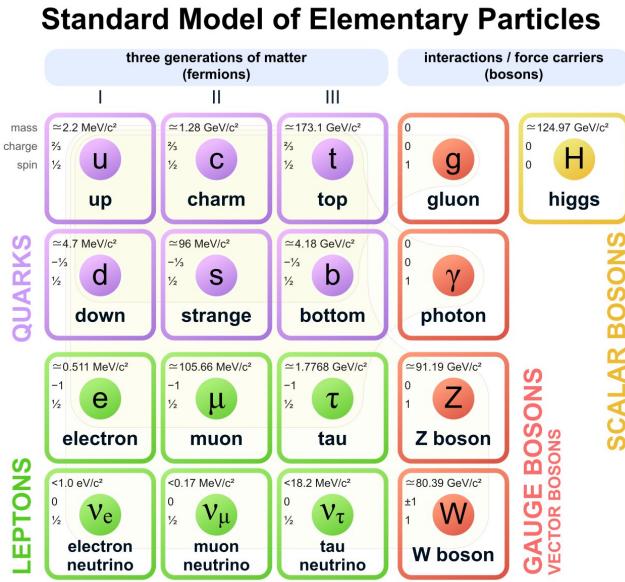


Figure 2.1: Overview of the SM of particle physics [17].

2.2.1 Fundamental interactions and gauge bosons

Spin-1 bosons have a very close connection to the fundamental forces of nature. Their quantum fields mediate the three fundamental forces of nature described by the SM: the *electromagnetic force*, the *weak force* and the *strong force*. Every single one of these forces is linked to a fundamental symmetry of nature and will interact with particles that contain a certain charge linked to that symmetry.

In the case of the **electromagnetic force**, this property is the well-known electric charge of particles. All particles that contain such a charge will feel each other's presence and interact accordingly. In the SM, the aforementioned $SU(2)_L \times U(1)_Y$ symmetry group breaks down to $U(1)_{EM}$ (see Section 2.3.2), which is responsible for the electromagnetic interaction. The theory behind this interaction is also referred to as quantum electrodynamics (QED). The gauge boson that mediates the electromagnetic force is the massless *photon*. Due to the Abelian nature of the $U(1)_{EM}$ group, it does not have an electric charge.

The strongest of the fundamental forces is appropriately called the **strong interaction**. It has a relative strength that is a factor 137 higher than the electromagnetic force and takes effect on all particles that contain a colour charge. While there is a single charge in QED, there are three charges in the case of colour charges: *red*, *green* and *blue*¹. The theory behind the strong interaction is called quantum chromodynamics (QCD) and is described by the non-Abelian group $SU(3)_C$. As such, 8 massless gauge bosons – called *gluons* – that interact with coloured particles roll out of the theoretical invariance requirements of this group. All gluons contain a

¹This naming framework is purely conventional and there is no link between the colour charge and what we perceive as a colour linked to the wavelength of light.

net colour charge themselves and are therefore allowed to self-couple to other gluons.

The fact that gluons contain a colour charge brings about an interesting concept called *colour confinement*. When we take two electromagnetic particles and start to separate them, the electric field between the particles will decrease as the field lines spread out. The opposite thing happens for two particles with a colour charge. When we start to separate them, the gluons will interact with each other and the field lines will not spread out. Instead, a tube of constant energy density is formed between the two coloured particles as illustrated in Figure 2.2. As the particles are separated, more and more energy needs to go into maintaining this constant flux tube until at a certain point there is so much energy that new coloured particles are produced to form pairs. This causes coloured particles to never exist in isolation but to bundle together into colourless compound particles.

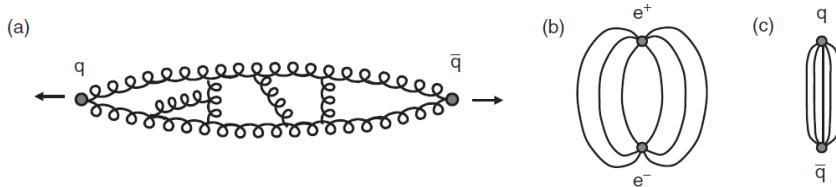


Figure 2.2: (a) Illustration of the self-interactions between the gluon fields between two quarks being separated. Electric field lines (b) are spread out and disperse at higher separation while strong field lines (c) form a tube of constant energy density [11].

Similarly, there is a related concept called *asymptotic freedom*. The “constant” that determines the strength of the coupling (called the coupling constant) does not remain constant but varies with the energy scale. When a charged particle is at a certain point in space, there will be unstable, short-lived particle-antiparticle pairs in the vacuum around it. For electromagnetic interactions, the further away two particles are from each other, the more these pairs “hide” the effective charge and the weaker the interaction becomes. The closer you get, the less of these pairs obscure the effective charge and the stronger the interaction gets. This effect is called “screening”. For the strong interaction, the effect is the opposite due to the colour charge of gluons. The smaller the length scale (or equivalently the higher the energy scale), the weaker the strong interaction gets.

The final force on the list of fundamental interactions is called the **weak force** with interactions that are described by the non-Abelian $SU(2)_L$ group. This means that similarly to gluons, the gauge bosons propagating the weak force are subject to self-interaction. The main difference with gluons is that the weak gauge bosons are massive.²

There are three gauge bosons that mediate the weak force. One of them is an electrically neutral boson called the *Z boson* clocking in at a mass of 91.19 GeV and the other two are

²The weakness of the weak interaction actually stems not from a low coupling constant, which has a value of $\frac{1}{30}$ [11], but from the mass of its gauge bosons. The propagators describing these interactions in quantum field theory are highly suppressed at low energies.

electrically charged states of a *W boson* (W^+ and W^-) with a mass of 80.38 GeV [18]. The first mediates neutral current interactions where the total charge in the initial and final states is electrically neutral. The charged *W* bosons allow for charged current interactions where the charge from the initial state can be transferred to the final state. A fascinating observation was made in 1956 by C.S. Wu and her collaborators when studying the radioactive decay of Cobalt 60 [19]. While the strong and electromagnetic interactions are symmetric under parity transformations (inverting the sign of the spatial coordinates), they noticed that parity is violated in the charged current weak interactions. Similar observations were made by Stanford Linear Accelerator Center (SLAC) in 1978 in the neutral current weak interactions [20].

To understand what this means we need to introduce the concepts of helicity and chirality. The first is defined as the relative alignment of the spin and momentum vectors of a particle. When the two vectors are aligned in the same direction, particles are called right-handed while spin and momentum vectors with opposite directions are referred to as left-handed. The chirality of a field is an internal quantum mechanical property related to projection operators. We will not go into the mathematical background here but it is important to know that in the massless limit, chirality and helicity become identical. Because of this, the names left-handed and right-handed are often adopted for chirality but it is important to remember that unlike helicity, chirality is an internal property that can not be changed by switching to a different frame of reference for massive particles. Parity conservation tells us that a force couples to left-handed and right-handed particles in the same way. Interestingly the abovementioned experiments found that the weak force violates this parity conservation and can distinguish between two particles with different chirality. In fact, the weak interaction is an extreme case of parity violation referred to as a *V-A* theory where there is no coupling at all between the weak force and right-handed particles or left-handed antiparticles [21, 22].

An honourable mention goes to the **gravitational force**, which is not included in the SM. While its effects are invaluable to our ability to not detach from the surface of our planet, it can currently not be described as a quantum theory. This is not an issue for present-day particle accelerator research as quantum gravity effects are expected to only become relevant at the Planck energy scale of order $\mathcal{O}(10^{16} \text{ TeV})$. The highest energies reached at accelerators are of the order of $\mathcal{O}(10 \text{ TeV})$, a scale where gravitational effects are negligible.

Another elementary boson that was not yet touched upon in this section is the **Brout-Englert-Higgs (BEH) boson** H . While the force-carrying bosons discussed before all have a spin of 1, the H boson is a scalar (spin-0) boson. Its field is responsible for giving mass to elementary particles through the mechanism of electroweak symmetry breaking (EWSB). Any particle that couples to the H field receives a mass because of it. We will divulge this mechanism in Section 2.3.

2.2.2 Fermions

The SM currently contains twelve spin- $\frac{1}{2}$ fermions that can be divided into two main categories.

Leptons

The six SM fermions without a colour charge are called the *leptons*. While they are all able to interact weakly, only three of them have an electric charge that is one negative fundamental charge unit in size. The lightest of these three is the world-renowned electron with a mass of 0.51 MeV, responsible for the electric current powering the laptop I am currently working on. The two other electrically charged leptons, namely the muon μ and the tau lepton τ are identical to the electron in every way except for their mass. A muon has a mass of 105.65 MeV, a little over 200 times larger than the electron mass, while the τ lepton tops out at an astounding mass of 1.78 GeV, making it the heaviest lepton [23]. These three variations of nearly identical particles are also called “generations” with the electron being the first generation, the muon the second generation and the tau the third generation. A single generation is also referred to as the flavour of the lepton (electron, muon or tau).

Every generation is linked to an electrically neutral lepton in an $SU(2)_L$ pair: electron neutrinos, muon neutrinos and tau neutrinos. Since these neutrinos do not have a colour charge or an electric charge they can only interact through the weak interaction, making interactions with neutrinos a rare occurrence. While their exact mass is still unknown, upper limits have been set at the eV level [24]. This means they have an incredibly low mass when compared to all other massive particles in the SM. We will come back to this topic in Section 2.4.2 and Chapter 3.

So far, we have described matter with their charge as we encounter them in our daily lives, such as the negatively charged electron. However, quantum field theory (QFT) also postulates the existence of *antiparticles* that are the same as the matter described so far in every aspect except for the sign of their charges. As such the negative electron has a positive “positron” counterpart. When a particle and antiparticle come together, they annihilate into energy and potentially different particle-antiparticle pairs.

A concept called *lepton number conservation* was brought into existence to explain the absence of a few expected processes in experiments [25, 26]. For every lepton flavour, a lepton flavour number is introduced with a value of 1 for leptons of the relevant family, -1 for antileptons of the relevant flavour and 0 for all other particles. The electron number L_e is for example 1 for electrons and -1 for positrons. Originally, lepton number conservation postulated that the sum of lepton numbers for each flavour separately should be conserved. A process where a particle with electron number 0 decays to 2 electrons would not be allowed. Over time we have found this number to only be approximately conserved as phenomena like neutrino oscillations (see Section 3.2.2) break lepton flavour number conservation. However, as far as we know, the total lepton number of all three families combined is a conserved quantity.

Quarks

The remaining six fermions that do have a colour charge are called the *quarks*. They follow a similar generation structure as the leptons, with every generation containing a quark with

positive electric charge $q = +\frac{2}{3}$ and a quark with negative electric charge $q = -\frac{1}{3}$. The three generations of quarks with positive electric charge are called up, charm and top quarks, which will be referred to as up-type quarks from now on. The other three quarks are called the down, strange, and bottom quarks and will similarly be referred to as down-type quarks. Similarly to the leptonic case, the higher the generation, the higher the quark mass.

Due to the principle of colour confinement, quarks can not exist as solitary particles. They are always bundled together into colourless combinations called *hadrons*. More specifically, hadrons made up of a quark-antiquark pair are called mesons while bound states of three quarks are called baryons. The most famous baryons are protons, consisting of two up quarks and a down quark, and neutrons, made up of two down quarks and an up quark.

In a similar fashion to lepton number conservation, a conservation law regarding baryons exists. A quantity called *baryon number* is defined as $\frac{1}{3}(n_q - n_{\bar{q}})$ with n_q the number of quarks and $n_{\bar{q}}$ the number of antiquarks in the bound state. The total baryon number sum in each process should be conserved.

2.3 Some concepts from Quantum Field Theory

In classical mechanics, a particle on a trajectory can be described by a Lagrangian L . The least action principle postulates that the trajectory that follows the classical equations of motion is the one where the integral of this Lagrangian over time, also known as the action, is minimized [27]. This concept can be expanded to a field theory where we no longer look at the description of a particle at a specific point in time and space but of fields that have values at every point in space-time. The Lagrangian is then replaced by a Lagrangian density \mathcal{L} and the action is defined as

$$S = \int d^4x \mathcal{L} . \quad (2.1)$$

The least action principle once again dictates that minimization of this action results in the field equations.

Quantum field theory marries the concepts of classical field theory, quantum mechanics and relativity. The underlying physics is also described by a Lagrangian density defined in space-time as a function of quantum fields and their derivatives. Following quantum mechanics, quantum fields are now defined as operators in each point in space as opposed to physical values. The technical details will not be covered here but in the spirit of quantum mechanics, fields can be quantized into an infinite sum of annihilation and production operators. While these fields are not physically measurable, particles come into existence as excitations of these fields [10, 28].

2.3.1 The standard model

The SM of particle physics is also described by a Lagrangian density. This density contains terms that describe free particles, interactions between particles and mass terms for the massive particles and is defined as:

$$\begin{aligned}
\mathcal{L} = & -\frac{1}{2}\text{Tr } G_{\mu\nu}G^{\mu\nu} - \frac{1}{2}\text{Tr } W_{\mu\nu}W^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} \\
& + (D_\mu\phi)^\dagger D^\mu\phi + \mu^2\phi^\dagger\phi - \frac{1}{2}\lambda(\phi^\dagger\phi)^2 \\
& + \sum_{f=1}^3 \left(\bar{l}_L^f i\cancel{D} l_L^f + \bar{l}_R^f i\cancel{D} l_R^f + \bar{q}_L^f i\cancel{D} q_L^f + \bar{u}_R^f i\cancel{D} u_R^f + \bar{d}_R^f i\cancel{D} d_R^f \right) \\
& - \sum_{f,g=1}^3 \left(y_l^{fg} \bar{l}_L^f \phi e_R^g + (y_l^{fg})^* \bar{e}_R^f \phi^\dagger l_L^f \right) \\
& - \sum_{f,g=1}^3 \left(y_d^{fg} \bar{q}_L^f \phi d_R^g + (y_d^{fg})^* \bar{d}_R^g \phi^\dagger q_L^f + y_u^{fg} \bar{q}_L^f \tilde{\phi} u_R^g + (y_u^{fg})^* \bar{u}_R^g \tilde{\phi}^\dagger q_L^f \right).
\end{aligned} \tag{2.2}$$

The different fields in this Lagrangian density are defined as follows: G is the gluon field of the strong force, W and B are the gauge fields of the $SU(2)_L$ and $U(1)_Y$ groups of electroweak interaction³, l are the lepton fields with e_R the right-handed charged leptons, q_L is the left-handed quark field while u_R and d_R are the right-handed up- and down type quarks respectively. Let's take a moment to dissect this Lagrangian density line by line.

The first line shows the gauge terms for the strong $SU(3)_C$ and the electroweak $SU(2)_L \times U(1)_Y$ theories. They mitigate the kinematics of the gauge fields and the self-coupling of the electroweak bosons and gluons. The second line represents the scalar H field, more specifically its kinematics and coupling to gauge bosons, its mass and its potential. The third line in turn is responsible for the kinematics of the leptons and quarks as well as the interaction of these particles to the gauge bosons. Both aspects are contained in the derivates \cancel{D} . For the term $\bar{l}_L^f i\cancel{D} l_L^f$ for example this derivative is defined as

$$\bar{l}_L^f i\cancel{D} l_L^f = \bar{l}_L^f i\gamma^\mu \left[\partial_\mu - ig_1 \frac{Y}{2} B_\mu + ig_2 \frac{\sigma^a}{2} W_\mu^a \right] l_L^f, \tag{2.3}$$

where γ^μ are the gamma matrices, Y is the weak hypercharge and σ^a are the three Pauli matrices. The first term describes the free kinematics of the leptons while the second and third term show coupling to the electroweak bosons (with g a coupling strength). Returning to the Lagrangian density in equation 2.2, the fourth term describes the coupling of leptons to the scalar H field and therefore the mass of these leptons. In a similar fashion, the last line shows the coupling of quarks to the H field. The summation in the last three lines runs over all three flavour generations.

³The W and B fields mix into the photons, W bosons and Z bosons of the electromagnetic and weak interactions discussed in the last section through the mechanism of EWSB.

2.3.2 Spontaneous symmetry breaking

Gluons, weak bosons and photons are all directly linked to one of the fundamental forces of nature but for the H boson, this is not the case. Why do we need this scalar boson in our model then? To answer this question we once again have to look at the importance of symmetry in physics. The fundamental interactions rising from the third line of the Lagrangian density in equation 2.2 are a direct consequence of the requirement that the Lagrangian should not only be globally invariant, but also locally invariant under gauge transformations at every point in space-time. The problem with this is that we know that W and Z bosons have a mass and the ad hoc inclusion of mass terms for these bosons to the Lagrangian would break local gauge invariance, leaving us with a contradictory and broken theory.

We are faced with an issue that was shipped in together with the parity violation in the weak sector. W bosons only interact with left-handed fermions (or right-handed anti-fermions) but do not “see” their right-handed partners whatsoever. If we want our fermions to have a mass however, we need a mass term for them in the Lagrangian of the form $m\bar{\psi}\psi$. One can also write it out in the following form:

$$m\bar{\psi}\psi = m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L) , \quad (2.4)$$

which will break gauge invariance under weak gauge transformations. To solve this issue while having masses for all the abovementioned particles, the concept of *spontaneous symmetry breaking* comes in [29–32].

Spontaneous symmetry breaking can be found in a plethora of systems such as ferromagnets and superconductivity. To understand the core concept behind it, we can look at an example of a marble being thrown in a bowl. The spherical symmetry of the bowl causes the ground state of the marble to be in the centre of the bowl, where it will be stable and thus will settle at the end of its trajectory. Imagine the centre of the bowl is now raised (like the bottom of a wine bottle) as illustrated in Figure 2.3. There is still a local extremum of the potential energy in the very centre of the bowl but it will be a maximum, meaning it is unstable and the smallest interaction will cause it to fall off. The ground state will now be settled at the ring at the bottom. However, there no longer is just one ground state but a multitude of ground states all set in that ring. Because one of these ground states has to be chosen, the symmetry is spontaneously broken. In other words, spontaneous symmetry breaking occurs when the ground state does not share the symmetry of the underlying physics. It is this exact idea that is behind the H field and how it gives mass to most particles through the mechanism of EWSB.

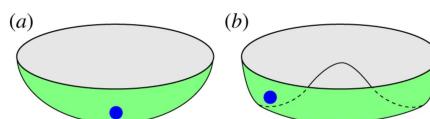


Figure 2.3: Illustration of an unbroken (a) and spontaneously broken (b) symmetry [32].

By adding a complex scalar $SU(2)$ doublet field ϕ of the form

$$\phi = \begin{bmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{bmatrix}, \quad (2.5)$$

this phenomenon can be introduced into the SM through the second term in 2.2. The potential $\mu^2\phi^\dagger\phi - \frac{1}{2}\lambda(\phi^\dagger\phi)^2$ has such a collection of ground states when μ^2 is negative. We can pick one by defining the ground state as:

$$\phi_{vac} = \begin{bmatrix} 0 \\ \frac{\mu}{\sqrt{2\lambda}} \end{bmatrix} = \begin{bmatrix} 0 \\ v \end{bmatrix}. \quad (2.6)$$

This can be interpreted as ϕ_3 in equation 2.5 being a real scalar field that causes fluctuations around v while ϕ_1, ϕ_2 and ϕ_4 are three “unphysical” fields (also called *Goldstone bosons*) that bring one ground state to a different ground state. By setting these last three to zero (called the unitary gauge) and inserting ϕ_{vac} into the first term on the second line of the Lagrangian density in equation 2.2, one can see mass terms for the weak bosons appear in the Lagrangian density. Because the three Goldstone bosons disappear and three massive gauge bosons appear along with a massive scalar boson, people also refer to this as the Goldstone bosons being absorbed by the gauge bosons.

In order to also get mass terms for the fermions into the Lagrangian density, we have to look at the interaction terms between the H field and the fermion fields. In equation 2.2 these are shown on the fourth and fifth lines. Let us take a closer look and choose ϕ as an expansion around the unitary gauge again: $\phi = [0 \ v + h]^T$ where h represents the real scalar field. The fourth line of 2.2 becomes:

$$\begin{aligned} \mathcal{L}_{h,lep} &= -\sum_{f=1}^3 y_l^f \left(\bar{l}_L^f \phi l_R^f + \bar{l}_R^f \phi^\dagger l_L^f \right) \\ &= -y_l^e \left((\bar{\nu}_e \ \bar{e}_L) \begin{bmatrix} 0 \\ v + h \end{bmatrix} e_R + \bar{e}_R [0 \ v + h] \begin{bmatrix} \nu_e \\ e_L \end{bmatrix} \right) + \text{muon and tau terms} \\ &= -y_l^e (v \bar{e}_L e_R + v \bar{e}_R e_L) + \text{interaction terms} + \text{muon and tau terms} \\ &= -y_l^e v \bar{e} e + \text{interaction terms} + \text{muon and tau terms}, \end{aligned} \quad (2.7)$$

which are the mass terms for the leptons in a gauge invariant setting. The mass of the fermions is directly proportional to the coupling of the fermion to the H field, also referred to as the Yukawa coupling. It can be seen that the H mechanism in itself does not generate masses for the neutrinos. These can be added through other methods which will be discussed in the next chapter. Similarly, the coupling terms for quarks to the H field determine the quark masses.

2.3.3 Feynman diagrams

When we get into the realm of calculating probability amplitudes, decay rates or scattering cross sections⁴ for interactions from the Lagrangian, the concept of perturbation theory kicks in. These calculations break down into a sum of contributions ever-increasing in the power (also called order) of a small-valued parameter, meaning that every higher-order term makes increasingly smaller contributions to the value of the total sum. In our case, this parameter is the coupling constant.

Every term in the perturbation theory can be associated with a set of Feynman diagrams. This is a mathematical tool that allows us to visualize all contributions and to break them down into components that can be digested more easily instead of simply having an “abstract” formula. An example for a common production mode in proton-proton colliders, called Drell–Yan (DY) production⁵, is given in Figure 2.4.

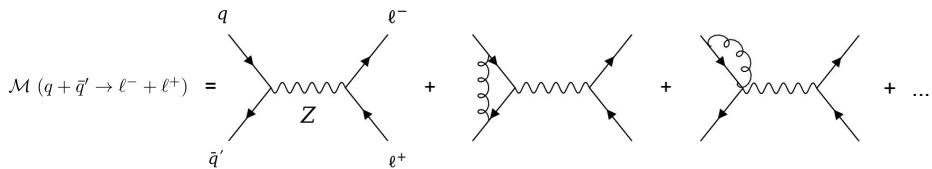


Figure 2.4: Feynman diagrams representing DY production.

Time is represented on the horizontal axis. One can see that for every vertex that is added, one goes higher in order. Besides being used for calculations in perturbation theory, Feynman diagrams are also used to visually represent a certain process, as is done in chapters 6 and 7. For this last purpose, we will always stick to the leading order (LO) diagrams for simplicity.

2.4 Limitations of the standard model

2.4.1 The Hierarchy Problem

EWSB in the SM predicts the existence of the scalar H boson and we have actually discovered this boson in the CMS and ATLAS experiments at the LHC in 2012. It truly was a great discovery, but also one that left us with a question we do not quite understand. This has everything to do with the mass of the scalar particle that was observed, which is settled at

⁴The cross section in particle physics terms is a measure of the probability that a specific process will take place.

⁵DY production refers to the resonant production of dilepton pairs.

125.25 GeV [23]. When taking a look at the one-loop quantum corrections to the H boson mass, they can be seen to add the following terms to the H boson mass squared due to a Dirac fermion [33]:

$$\Delta M_H^2 = N_f \frac{\lambda_f^2}{8\pi^2} \left[-\Lambda^2 + 6m_f^2 \log \frac{\Lambda}{m_f} - 2m_f^2 \right] + \mathcal{O}\left(\frac{1}{\Lambda^2}\right). \quad (2.8)$$

Here m_f is the mass of the fermion in the loop, N_f is the number of fermions with mass m_f , λ_f is the Yukawa coupling to the fermion and Λ is the so-called cut-off scale. This last variable is the energy scale at which the QFT breaks down and new physics takes over. We can set it for example at the Planck scale where quantum gravitational effects grow strong ($\mathcal{O}(10^{19} \text{ GeV})$). From the quadratically divergent behaviour observed in equation 2.8, one would in that case expect the H boson mass to be very close to the Planck scale and not at the value it is observed at. In principle one could get away with careful fine-tuning on a very precise level, but this would be very unnatural.

2.4.2 Neutrino mass

It was concluded earlier that EWSB does not generate masses for the neutrinos in the absence of right-handed neutrinos. Observations of neutrino oscillations, which we will return to in the next chapter when we discuss neutrino physics and related extensions to the SM, show us that at least two of the neutrinos in fact have to be massive particles. This means that we need mass terms for these particles in the Lagrangian density. In principle it is possible to add such gauge invariant mass terms by introducing right-handed Dirac neutrinos but they would be directly proportional to the H vacuum expectation value v [34]. Because there are upper limits on the neutrino masses at the eV level [35–37], this would pose a new uncomfortable situation in the form of a hierarchy problem. Why is the neutrino mass so much lower than that of all other massive SM particles? This question is one of the key motivations for the research presented in this dissertation. One possible answer is to introduce additional heavy neutrinos in the model, as will be explained in the next chapter. These additional heavy neutrinos will be the main topic of the search in Chapter 7.

2.4.3 Dark matter

Experiments in the fields of cosmology and astrophysics have observed a number of results that point towards the existence of dark matter. This is matter that does not seem to interact through the electromagnetic force and therefore can not be seen with light. One of these observations is the velocity of rotation of spiral galaxies as a function of radial distance [38]. Most of the mass in these spiral galaxies is contained in the centre of the galaxy which would lead one to expect that the rotation velocity will decrease with radial distance as per Kepler's second law. What is observed, however, is that the velocity remains stagnant as can be seen in Figure 2.5. This could be explained by the existence of a dark matter halo around the galactic centre.

There is a wide range of other results to back up the existence of dark matter such as

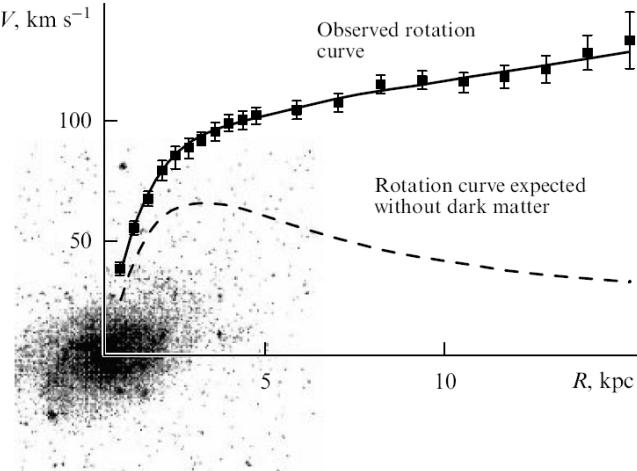


Figure 2.5: Observed rotation curve of the M33 galaxy versus the curve expected without dark matter [38].

fluctuations in the cosmic microwave background and structure formation after the Big Bang [39, 40]. While there is an argument for a modified theory of relativity to explain some of these observations, there are other observations that are much harder to explain by such a modified theory, one of which is the Bullet Cluster [41]. This is a cluster that is being formed from the collision of two predecessor clusters. X-ray images show that the visible matter is not in the same location as the bulk of the mass observed through gravitational lensing. It can be seen in Figure 2.6 that the visible mass lags behind the bulk of the mass, which can easily be explained by dark matter but is not as straightforward with modified relativity theories.

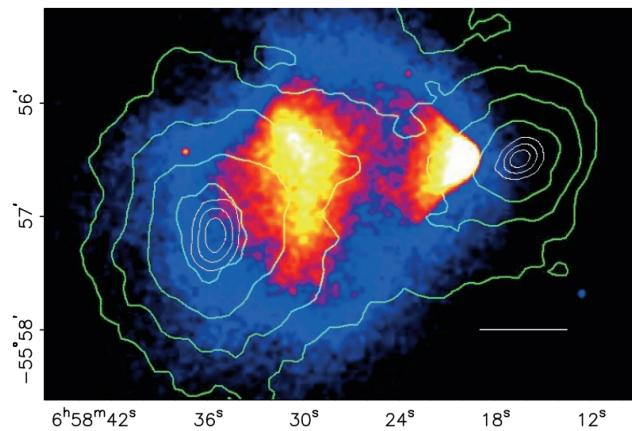


Figure 2.6: X-ray images of the bullet cluster. The green contours show the mass distribution obtained through gravitational lensing observations [41].

As it stands, the SM does not offer a candidate for dark matter. Most particles in the SM are electrically charged and are thus immediately disregarded in this context. The only stable, electrically neutral particles that are contained in the SM are the neutrinos. Here we run into a cosmological wall however. As neutrinos are created ultra-relativistically at the earlier stages

of the universe on account of having a low mass, they do not allow for the structure formation we observe nowadays [42]. The neutrinos would simply stream freely. Therefore, if dark matter is a type of particle, it must be one that we currently do not have included in our SM.

2.4.4 Matter-antimatter asymmetry

Taking a look at the overview of the SM content immediately raises one question. Where is all the antimatter? In our daily lives, we mostly encounter matter but from cosmological models, one would expect them to be created in equal amounts just after the Big Bang [43, 44]. Experiments like the Alpha Magnetic Spectrometer (AMS) detector are looking for signatures of antimatter stars by looking for the presence of anti-helium but so far nothing has turned up [45]. There are mechanisms in the SM that allow for so-called CP violation [46] – meaning charge-parity symmetry violation – to end up with slightly more matter than antimatter but it is far from sufficient to explain the large difference we observe in the universe. To explain the matter-antimatter asymmetry question at the level of differences we observe, additional sources of CP violation that are currently not part of the SM should be included.

3

Extending the standard model

It is clear from the last chapter that the SM has not yet reached its final form. Many open questions have either an unsatisfying solution in the SM or no solution at all. In order to address these questions, physicists have constructed numerous models that expand or adapt the SM. Two of them form the subject of searches described in later chapters where the goal was to either discover the presence of new particles predicted by the models or exclude them as best as we can. We will take some time to highlight them in this chapter.

We will start with an overview of SUSY [33, 47, 48] in Section 3.1, followed by a description of the specific type of model that is relevant to the search in Chapter 6. Because it can offer an explanation to a wide range of questions like the hierarchy problem and the origin of dark matter, SUSY has grown into a popular theory among particle physicists. The first of these, the hierarchy problem, will be used to introduce the concept of SUSY in Section 3.1.1.

Section 3.2 will then move on to delve into extensions of the SM in the neutrino sector. The first subsection expands on the issues in this sector that were touched upon in the last chapter. This will be followed by an introduction to the seesaw mechanism along with the necessary theory and models to prepare us for the search for heavy neutral leptons in Chapter 7. Finally, an overview of the current state-of-the-art of such searches is given to close out the chapter.

3.1 Supersymmetry

3.1.1 Solving the hierarchy problem

Equation 2.8 showed us that the contribution of the one-loop corrections due to fermions to the H boson mass squared is quadratically divergent. We can do the same exercise for one-loop corrections due to a scalar particle. This correction takes the following form [33]:

$$\Delta M_H^2 = N_S \frac{\lambda_S}{16\pi^2} \left[-\Lambda^2 + 2m_S^2 \log \frac{\Lambda}{m_S} \right] - \frac{\lambda_S^2 N_S}{16\pi^2} v^2 \left[-1 + 2 \log \frac{\Lambda}{m_S} \right] + \mathcal{O}\left(\frac{1}{\Lambda^2}\right), \quad (3.1)$$

where this time m_S is the mass of the scalar, N_S is the number of scalars with mass m_S , λ_S is the coupling between the H boson and the scalar and v is the vacuum expectation

value of the H field. This correction has a familiar shape to that of the fermionic one-loop corrections. In fact, if one assumes $N_S = 2N_f$ and $\lambda_f^2 = -\lambda_S = 2m_f^2/v^2$, the fermion and scalar contributions to the H boson mass become:

$$\Delta M_H^2 = N_f \frac{\lambda_f^2}{4\pi^2} \left[(m_f^2 - m_S^2) \log \frac{\Lambda}{m_S} + 3m_f^2 \log \frac{m_S}{m_f} \right] + \mathcal{O}(\frac{1}{\Lambda^2}) \quad (3.2)$$

and the quadratic divergence is gone. There is still a logarithmic divergence present in equation 3.2 but when $m_S = m_f$, even that divergence disappears. This means that when the following conditions are met:

1. The number of scalars with mass m_S is twice the amount of number of fermions with mass m_f
2. The bosonic and fermionic couplings are related in a scale-invariant way

the quadratic divergences disappear. This can be achieved by a symmetry that links bosons and fermions as well as their couplings. The idea of cancellation of quadratic divergences can be also propagated to the other particles in the SM such as the W and Z bosons by introducing fermions with couplings to the H boson that are linked to those of the weak bosons.

A glance at the SM content shows that the two conditions above are not met within the SM. Therefore we would have to introduce a new symmetry that brings along additional fermions for every boson and vice versa. This is exactly what SUSY does. Ideally, these extra fermions and bosons would have the same mass as their SM counterparts. From experiments, however, we know that this is not the case or we would have already observed these particles. That means that SUSY would need to be softly broken, a concept we will return to later on. In order to avoid another hierarchy problem with the logarithmically divergent terms, the masses of these new particles are expected to be below the order of 1 TeV.

3.1.2 Supersymmetry and soft breaking

The idea behind supersymmetry is to have a symmetry that relates bosons and fermions. SUSY generators \mathcal{Q} transform one into the other:

$$\mathcal{Q} |\text{fermion}\rangle = |\text{boson}\rangle \quad , \quad \mathcal{Q} |\text{boson}\rangle = |\text{fermion}\rangle . \quad (3.3)$$

The SM particles and their superpartners (sparticles) can be combined into superfield multiplets that contain equal fermionic and bosonic degrees of freedom. There are many candidate symmetry groups which with one can construct a supersymmetric model but as discussed in the last section they all should be broken at the energy scale of the electroweak theory. As the masses of the new sparticles can not be too high, since that would introduce a new fine-tuning problem, we call this soft breaking of SUSY.

Up to now, no fully satisfactory solution has been found to fundamentally break SUSY. Instead, it is often chosen to work with effective field supersymmetric theories where the SUSY breaking terms are added to the Lagrangian density by hand [49, 50]. The concept

behind an effective field theory is that the actual breaking mechanism happens in a “hidden” sector that is at an energy scale higher than the soft SUSY scale, which is potentially out of reach in this day and age. This would be a sort of black box that causes the effects that warrant the existence of the breaking terms in the Lagrangian density at the electroweak energy level. The same idea can be translated to e.g. nuclear physics, which can be interpreted as an effective theory of the SM interactions at lower energies. One example of such an effective SUSY theory that we will go into in the next subsection is called the *Minimal Supersymmetric Standard Model (MSSM)* [51].

3.1.3 Minimal Supersymmetric Standard Model

The MSSM is an effective supersymmetric model that aims to introduce SUSY with the least number of additions to the SM. The number of generators \mathcal{Q} in the MSSM is one. An overview of the particle and gauge content is provided below.

Matter content

As the changes to the SM are kept minimal, the MSSM contains the same three generations of quarks and leptons, which are all supplemented with spin-0 supersymmetric partners. As all fermions come in left-handed and right-handed chirality – except for the neutrinos –, left- and right-handed superpartners have to be added to maintain the same number of degrees of freedom. Notice, however, that since these sparticles are scalar and thus have no spin, the naming left- and right-handed have nothing to do with their chirality. Instead, it simply refers to the chirality of their SM superpartner. The according superfields Q, U_R, D_R, L, E_R are made up of the left- and right-handed SM fields along with their superpartners. An overview of these superfields and their matter content is shown in Table 3.1

Table 3.1: The superfields of the MSSM corresponding to the SM fermion content. Each row shows a different superfield and its contents.

| Names | | spin-0 | spin- $\frac{1}{2}$ |
|---|-------|-----------------------------|---------------------|
| squarks, quarks ($\times 3$ families) | Q | $(\tilde{u}_L \tilde{d}_L)$ | $(u_L d_L)$ |
| | U_R | \tilde{u}_R | u_R |
| | D_R | \tilde{d}_R | d_R |
| sleptons, leptons ($\times 3$ families) | L | $(\tilde{\nu} \tilde{e}_L)$ | (νe_L) |
| | E_R | \tilde{e}_R | e_R |

Gauge and Higgs content

All spin-1 gauge bosons in the SM are supplemented by a fermionic spin- $\frac{1}{2}$ superpartner in the MSSM. These superpartners are referred to as gauginos. The eight gluons have eight superpartners called gluinos (\tilde{g}) while the electroweak sector receives four extra fermions. In the BEH sector, it is important to note that in order to be able to give mass to both up-type and down-type quarks in the MSSM and to avoid chiral anomalies, a second BEH doublet

H_2 with opposite weak hypercharge to the already described BEH doublet H_1 needs to be introduced [33]. This means that after EWSB, there are five BEH particles: two neutral bosons h and H , a pseudoscalar A boson and two charged H^\pm bosons. The two original BEH doublets each have two fermionic higgsino superpartners, one of which is electrically neutral and one of which is electrically charged.

These four superpartners of BEH sector will mix with the electroweak superpartners to form mass eigenstates. As such we end up with four neutral fermions called neutralinos $\tilde{\chi}_i^0$ and four electrically charged fermions called charginos $\tilde{\chi}_i^\pm$, two with a positive charge and two with a negative charge. The subscript i ranks charginos and neutralinos according to their mass. A chargino with subscript 1 will have lower mass than a chargino with subscript 2. An overview of the charginos and neutralinos in the MSSM is given in Table 3.2. Chapter 6 describes a search for the production of the lightest charginos and neutralinos in the CMS detector.

Table 3.2: Superpartners of the SM boson sector.

| Superpartner names | SM boson | spin- $\frac{1}{2}$ superpartner |
|----------------------|--------------------------------|--|
| gluino | g | \tilde{g} |
| chargino, neutralino | W^\pm, Z, h H, A, H^\pm | $\tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0$ $\tilde{\chi}_1^+, \tilde{\chi}_2^+, \tilde{\chi}_1^-, \tilde{\chi}_2^-$ |

3.1.4 R-parity

One aspect of SUSY that was not touched upon so far is conservation of lepton and baryon numbers. In the MSSM, these numbers are not necessarily conserved. This has a rather problematic consequence for the stability of the protons, which now have the option to decay through the process $p \rightarrow e^+ + \pi^0$ [48]. Experimental evidence shows that protons are stable particles with lifetimes way beyond the lifetime of the universe [23], however.

A solution to maintain the stability of the proton in supersymmetric models is to introduce a new conservation law in the form of a new multiplicative quantum number called R-parity. It is defined as:

$$P_R = (-1)^{2s+3B+L}, \quad (3.4)$$

where s is the spin of the particle, B is the baryon number of the particle and L is the lepton number of the particle. Particles contained in the SM will have R-parity number +1 while their supersymmetric partners will have R-parity number -1. The product of this number for all particles in a process should be conserved, which has the important implication that a sparticle can not decay to a fully SM final state as it is impossible to get -1 from the multiplication of $+1 \cdot n$. Likewise, sparticles can only be produced from SM particles in pairs.

Conservation of R-parity actually supplies us with a great opportunity to define dark matter. Since the lightest sparticle will be massive and stable, it is the perfect candidate for this elusive matter type if it is not charged.

3.1.5 Supersymmetry at the LHC

The LHC experiments have a very broad programme for searches of SUSY signatures that aim to scan a large range of possible SUSY configurations. We will take a moment now to look a bit more closely at this landscape here. It should be noted that due to the very high number of free parameters in the SUSY models, it would be nearly impossible to perform a search and interpretation where all of the parameters are kept free. Instead one often makes use of simplified benchmark SUSY models that describe a subset of predicted sparticles, their production and their decay based on phenomenology [52–54].

The main production processes for squarks, gluinos, sleptons and gauginos at hadron colliders are shown in Table 3.3. It can be seen that there are two main types of sparticle production at hadron colliders, namely strong production and electroweak production. The cross sections for both types at the LHC are shown in Figure 3.1 as a function of sparticle mass.

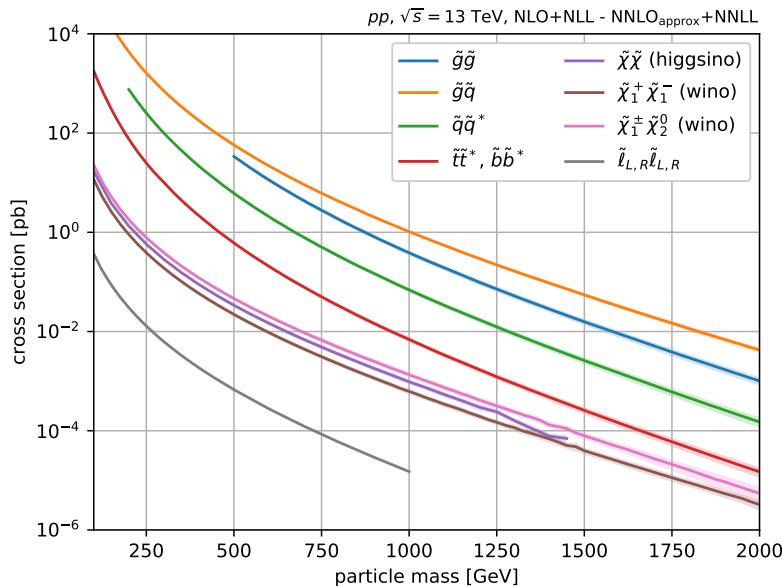


Figure 3.1: Production cross sections for different superpartner pairs in simplified MSSM topologies [56].

Table 3.3: Main production channels for sparticles in hadron colliders [55].

| Interaction | Process |
|-------------|---|
| Electroweak | $q\bar{q} \rightarrow \tilde{\chi}_i^+ \tilde{\chi}_j^-, \tilde{\chi}_i^0 \tilde{\chi}_j^0$ |
| | $q\bar{q} \rightarrow \tilde{l}_i^+ \tilde{l}_i^-, \tilde{v}_1 \tilde{v}_1^*$ |
| Strong | $u\bar{d} \rightarrow \tilde{\chi}_i^+ \tilde{\chi}_j^0$ |
| | $u\bar{d} \rightarrow \tilde{l}_L^+ \tilde{v}_1$ |
| | $d\bar{u} \rightarrow \tilde{\chi}_j^- \tilde{\chi}_i^0$ |
| | $d\bar{u} \rightarrow \tilde{l}_L^- \tilde{v}_1^*$ |
| | $gg \rightarrow \tilde{g}\tilde{g}, \tilde{q}_i \tilde{q}_j^*$ |
| | $gq \rightarrow \tilde{g}\tilde{q}_i$ |
| | $q\bar{q} \rightarrow \tilde{g}\tilde{g}, \tilde{q}_i \tilde{q}_j^*$ |
| | $qq \rightarrow \tilde{g}\tilde{g}, \tilde{q}_i \tilde{q}_j$ |

Squark and gluino production occurs through the strong interaction process between quark and/or gluons. As expected, the production cross sections for these interactions are also the highest at the LHC. This makes the strongly produced SUSY processes exquisite probes to perform searches for SUSY at the LHC experiments. The exact decay channels of the squarks and gluinos will highly depend on the model parameters and sparticle masses. So far no sign of SUSY in these production modes has been observed in any of the LHC experiments. Instead, exclusion limits on the masses of these hypothetical sparticles have been derived.

The latest summary plot from CMS on gluino pair production is shown in Figure 3.2. It shows searches in multiple decay modes to quarks and the lightest supersymmetric particle (LSP) that exclude gluino masses around 2 TeV. ATLAS has performed similar searches and places exclusion limits on gluino masses up to 2 TeV as well [57]. A similar summary plot from CMS on squark pair production can be found in Figure 3.3. It shows searches for scalar stop pair or b squark pair production in various decay channels that provide exclusion limits on their masses up to 1.25 TeV. Searches for light squark production show exclusion limits on their mass up to 1.75 TeV. Results from ATLAS on the same production channels show comparable limits [57].

The fact that these excluded masses already go well into the TeV range can have multiple implications. The mass of the colour sparticles might be too high to detect at the LHC or the mass spectrum could be too compressed - i.e. the masses of the produced sparticles are very close to one another - and the final state particles are very low in energy. In that case, it might be more advantageous to probe **electroweak production channels**. While their processes have lower cross sections, their models are generally less constrained. The latest summary plot of the latest CMS results, which was produced in 2021, can be found in Figure 3.4 and shows exclusion limits on the mass of charginos and neutralinos in simplified MSSM up to 1.3 TeV and on the mass of sleptons up to 700 GeV. The search in Chapter 6 falls in the category of electroweak sparticle production and targets general chargino-neutralino production in multilepton final states at CMS.

Lastly, it should be mentioned that while we focussed on MSSM-type simplified models, both ATLAS and CMS perform searches for different models and exotic signatures. As they fall outside the scope of this thesis, they will not be discussed here and the interested reader is referred to the public results pages of both experiments.^{1,2}

¹<https://cms-results.web.cern.ch/cms-results/public-results/publications/SUS/index.html>

²<https://twiki.cern.ch/twiki/bin/view/AtlasPublic/SupersymmetryPublicResults>

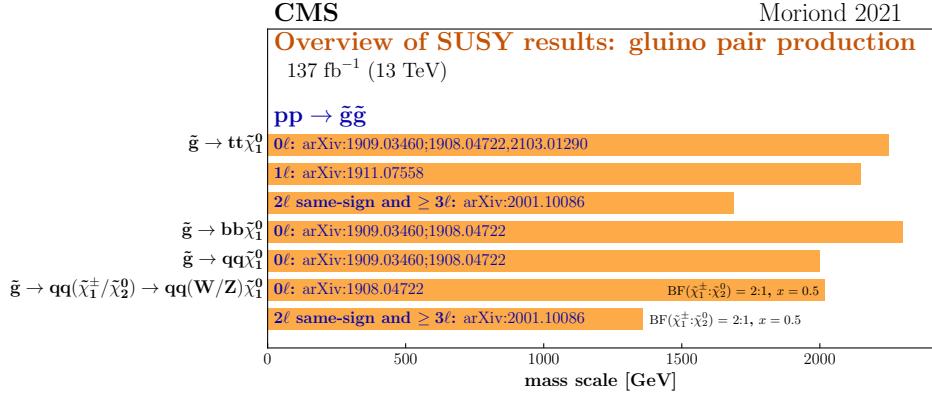


Figure 3.2: Latest summary plot from CMS showing the mass reach of Run II CMS SUSY results for processes with gluino pair production [58].

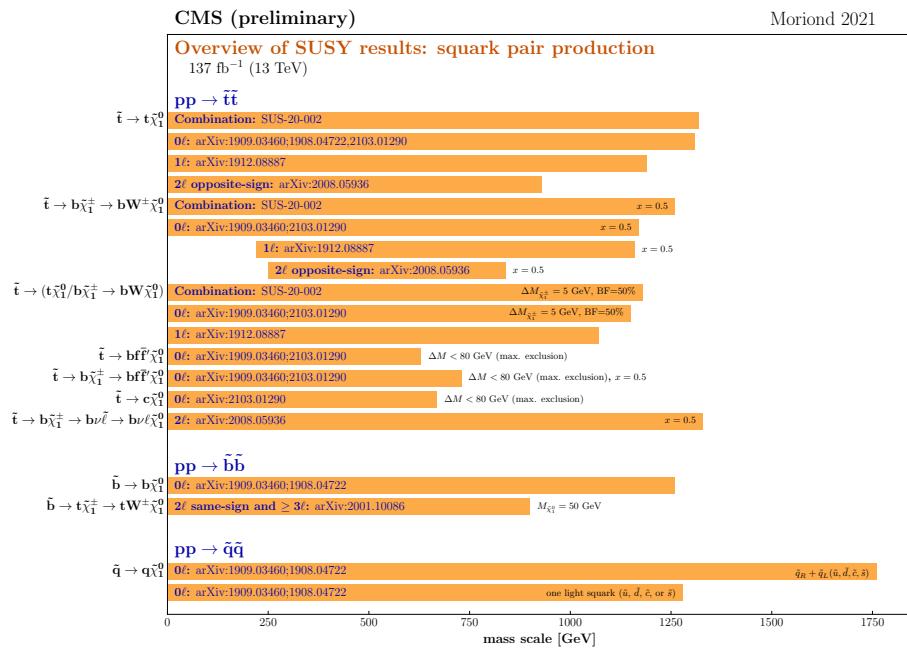


Figure 3.3: Latest summary plot from CMS showing the mass reach of Run II CMS SUSY results for processes with squark pair production [58].

3.2 Neutrino Physics

3.2.1 Introduction

In Chapter 2 the smallness of the neutrino masses was listed as one of the open questions of the SM. This leaves one to wonder why it has to be a problem, can neutrinos not just be massless particles? The answer to this by experimental observations in the form of neutrino oscillations is a resounding *no*. In the next subsection, we will go over the concept of neutrino oscillations and why it means that at least two of the SM neutrinos must have a nonzero mass

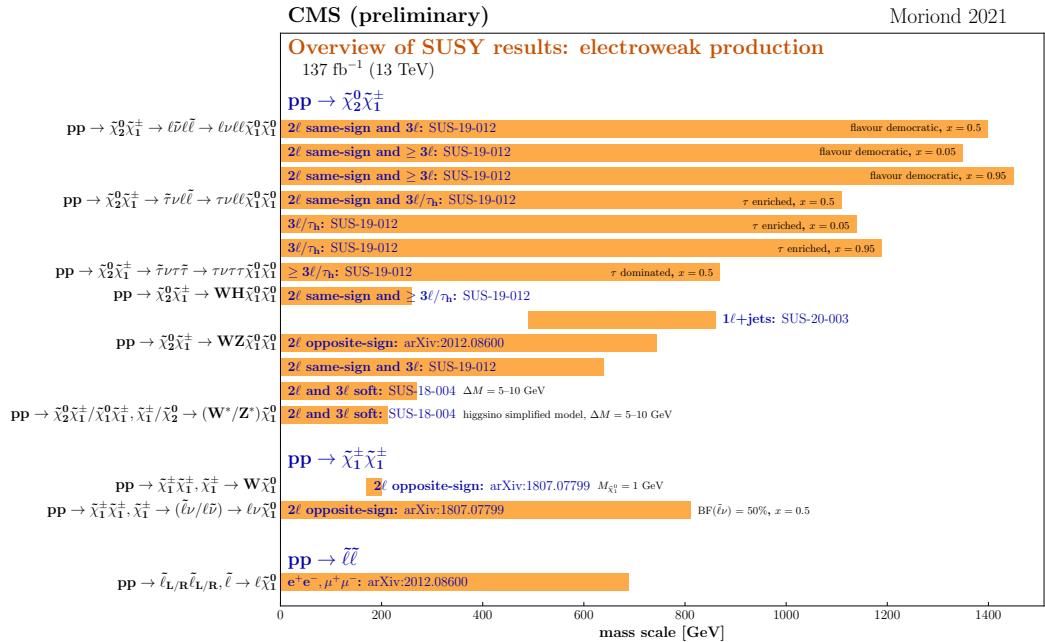


Figure 3.4: Latest summary plot from CMS showing the mass reach of Run II CMS SUSY results for electroweak production processes [58].

as well as the experimental observation of this phenomenon. Afterwards, we will explore different hypothetical mechanisms to give neutrinos their small mass values and set them up in light of the search for new physics in the neutrino sector in Chapter 7.

3.2.2 Neutrino Oscillations

Neutrinos have been defined in their flavour base up to now. Electron, muon and tau neutrinos are flavour eigenstates corresponding to the flavour of their charged lepton counterpart. There is no reason however that the flavour eigenstates have to correspond to the mass eigenstates that have a definite mass. In fact, one can define a mixing matrix that can be used to translate the flavour eigenstates into superpositions of the mass eigenstates and vice versa.

Things get interesting when we start propagating these superpositions through space. Imagine an electron neutrino is produced at a certain source. As this neutrino propagates, the different masses of the different mass eigenstates will cause them to propagate with a different phase. This difference in phase causes the neutrino to no longer be the superposition of mass eigenstates that defines an electron neutrino but could be a superposition of different flavour eigenstates. An electron neutrino that propagates through space could very well be measured at a different point on its trajectory as a muon neutrino! In the case of two flavours the probability of measuring a different flavour neutrino is given by:

$$P = \sin^2 2\theta \sin^2 \left(1.27 \Delta m^2 \frac{L}{E} \right), \quad (3.5)$$

where L is the distance travelled between the source and the detector expressed in km while E

is the energy of the neutrino in GeV. The mixing angle, denoted by θ , defines the amplitude of the neutrino oscillations. If it is 0 or $\frac{\pi}{2}$, it would mean there is no oscillation between the two involved states. $\Delta m^2 = m_2^2 - m_1^2$ in turn is the squared mass difference between the two mass eigenstates and it is exactly this dependency that forces us to the conclusion that at least two neutrinos must have mass. If two of the three neutrinos are massless, their squared mass difference would be zero and there would be no oscillation between the two states.

3.2.3 Dirac, Majorana and Weyl fermions

There are three types of fermions that could hypothetically exist, namely Dirac fermions, Weyl fermions and Majorana fermions [59].

Lets us start with the **Dirac fermions**, as all charged fields in equation 2.2 are of such type. They obey the well-known Dirac equation [11]:

$$(i\gamma^\mu \partial_\mu - m)\Psi = 0 , \quad (3.6)$$

where Ψ is the Dirac field, which is described by four-component Dirac spinors carrying four degrees of freedom. Depending on the representation one is working in, these components are open for interpretation. There is for example the chiral representation of these spinors that provide a description in terms of chirality. In the popular Dirac representation [60], the four components describe particle-antiparticle in spin-up and spin-down states. Important to note is that Dirac fermions have a clear distinction between particles and antiparticles. They are different objects with opposite charges and identical masses.

When fermions are massless, they can be described by the simpler **Weyl** equation [59]:

$$i\gamma^\mu \partial_\mu \psi = 0 . \quad (3.7)$$

This time the solutions can be described by two-component spinors or so-called Weyl spinors. As the mass of Weyl fermions is zero, it is not possible to boost into a frame where the helicity changes sign. The helicity is therefore unambiguously positive or negative and equals the chirality of the fermion; they are either left-handed or right-handed. It is possible to construct Dirac fermions from two Weyl spinors: one left-handed and one right-handed spinor.

Lastly, we come to the **Majorana fermions**. They are a construct by Ettore Majorana to see if neutrino mass terms could be written down without the need for right-handed neutrinos [61]. He succeeded in finding these terms by defining a type of fermion where the particle is its own antiparticle. The implication of this is that Majorana fermions need to have zero charges. As most fermions in the SM have some type of gauge charge, this means only neutrinos are potential candidates to be Majorana fermions after EWSB. This type of fermion is characterized by a four-component spinor where only two components are independent as the fermion and anti-fermion components are equivalent. While Majorana and Weyl spinors are not equivalent, it is possible to build a four-component Majorana spinor from a two-component Weyl spinor. Notice that since Dirac fermions can be built from two Weyl fermions, they can

also be built from two degenerate – i.e. having the same mass – Majorana fermions. A Dirac fermion is thus equivalent to two degenerate Majorana fermions [62].

3.2.4 Massive neutrinos

Of the two types described in the last subsection, neutrinos could theoretically be either Majorana or Dirac fermions. This means that additional neutrinos can get their mass from either Dirac mass terms, Majorana mass terms or both. Let us take a closer look at the implications of both types [62].

Dirac mass terms are defined by the existence of a left-handed and right-handed particle and are of the form:

$$m\bar{\psi}\psi = m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L) . \quad (3.8)$$

It can be interpreted as the coupling of the process where a left-handed particle is created and a right-handed particle is destroyed and vice versa. They are the type of mass terms that were discussed in Chapter 2 when dealing with spontaneous symmetry breaking. Section 2.4.2 showed that the absence of right-handed neutrinos causes the BEH mechanism to not generate neutrino masses. One can, however, manually add a set of three right-handed neutrinos to the SM that are singlets under every gauge group to re-introduce such mass terms. Taking another look at the Lagrangian density with these added right-handed neutrinos shows us:

$$\mathcal{L}_m^{\text{Dirac}} = - \sum_{f,f'=1}^3 y_l^{ff'} \bar{l}_L^f \phi \nu_R^{f'} + h.c. \quad (3.9)$$

for which the relevant neutrino mass terms after symmetry breaking are

$$\mathcal{L}_m^{\text{Dirac}} = - \frac{v}{\sqrt{2}} \sum_{f,f'=1}^3 \bar{\nu}_L^f y_l^{ff'} \nu_R^{f'} + h.c. \quad (3.10)$$

so that the neutrino masses after diagonalization of $y_l^{ff'}$ are

$$m_\nu = \frac{v}{\sqrt{2}} y_l^f . \quad (3.11)$$

With v being the BEH vacuum expectation value, this means that the smallness of the neutrino masses would have to be explained by the Yukawa coupling y_l^f , which raises the question of why it is so much smaller than that of the coupling of other fermions to the H boson. Thus a new hierarchy problem is born.

Since neutrinos carry zero electric charge, it is not necessary for right-handed neutrinos to exist in order to define **Majorana mass terms**, as Ettore Majorana showed. One could look towards the charge-conjugated left-handed field to build a mass term from purely left-handed fields. The charge conjugation operation is defined as the operation that inverts all charges of

a particle, turning it into its antiparticle. Majorana mass terms can then be written down as follows:

$$\mathcal{L}_m^{\text{Majorana}} = -\frac{1}{2} \sum_{f=1}^3 m_\nu \bar{\nu}_L^c \nu_L + h.c. . \quad (3.12)$$

While we will not go into detail here, it can be shown that it can be obtained from a mass term before symmetry breaking that is a dimension 5-point interaction between the neutrino and the H field. The mass m_ν will then be of the form αv^2 where α is inversely proportional to a new physics scale [63]. In contrast to the Dirac fermions, this proportionality does naturally explain the large difference in neutrino mass and the mass of all other particles. A major issue here is that this term will break $SU(2)_L \times U(1)_Y$ symmetry, an issue that can be resolved by adding heavy right-handed neutrinos as will be discussed in the next section.

Before moving to this discussion, let us close the subsection with two final observations on Majorana mass terms. Firstly, the mass term can be interpreted as the process that creates two neutrinos and will thus break lepton number conservation. Secondly, it is important to note that because Dirac neutrinos can be written as a combination of two degenerate Majorana neutrinos, it is possible for Dirac neutrinos to have two Majorana mass terms: one left-handed Majorana neutrino and one right-handed Majorana neutrino. Similarly, two degenerate Majorana neutrinos can be described by a Dirac mass term.

3.2.5 The Seesaw Mechanism

To solve the electroweak symmetry violation issue from the last subsection, one needs to introduce new, heavy particles. One way to achieve this is through the so-called *seesaw mechanism*. There are commonly three types of seesaw models. The first, appropriately called the Type-I seesaw mechanism, introduces heavy, right-handed singlet neutrinos, also referred to as heavy neutral leptons (HNLs), into the model. It is the mechanism of interest for the search in Chapter 7 and will be the focus of the next sections. A second type, called Type-II, adds a scalar triplet to the model, which extends the BEH sector of the SM [64, 65]. One of the new H bosons is doubly charged, which can decay to same-sign decay products that create an interesting signature in collider experiments [66–68]. A third type of seesaw mechanism, called type-III, adds an additional fermion triplet to the SM [69]. This expresses itself as the addition of two charged leptons and a neutral lepton to the model, which can be looked for in collider searches [70]. We will not discuss these last two seesaw models any further here and instead turn our attention to the Type-I seesaw model. Let us start by writing down the possible mass terms we can have after introducing singlet neutrinos [71].

We could have a Majorana mass term for the left-handed neutrino:

$$\mathcal{L}_\nu^{LL} = -\frac{1}{2} m_\nu^{LL} \bar{\nu}_L^c \nu_L + h.c. \quad (3.13)$$

But in order not to violate electroweak symmetry, m_ν^{LL} should be equal to 0. Similarly, we

can have a Majorana mass term for the right-handed neutrino:

$$\mathcal{L}_\nu^{RR} = -\frac{1}{2}m_\nu^{RR}\bar{\nu}_R\nu_R^c + h.c. \quad (3.14)$$

Since these right-handed neutrinos are sterile, meaning they are neutral with respect to all SM gauge interactions, this mass term does not cause any issues and the mass m_ν^{RR} does not have to be zero. Finally, we can also add the Dirac mass terms for the left- and right-handed neutrinos as well as the charge conjugates:

$$\mathcal{L}_\nu^{LR} = -m_\nu^D\bar{\nu}_R\nu_L - m_\nu^D\bar{\nu}_L\nu_R^c + h.c. , \quad (3.15)$$

where m_ν^D is the Dirac mass, which has to be equal for both terms. Combining all three mass terms gives us:

$$\begin{aligned} \mathcal{L}_\nu^m &= \mathcal{L}_\nu^{LL} + \mathcal{L}_\nu^{RR} + \mathcal{L}_\nu^{LR} \\ &= -\frac{1}{2}m_\nu^{LL}\bar{\nu}_L\nu_L^c - \frac{1}{2}m_\nu^{RR}\bar{\nu}_R\nu_R^c - m_\nu^D\bar{\nu}_R\nu_L - m_\nu^D\bar{\nu}_L\nu_R^c + h.c. \\ &= -(\bar{\nu}_L^c \quad \bar{\nu}_R) \begin{pmatrix} 0 & m_\nu^D \\ m_\nu^D & \frac{1}{2}m_\nu^{RR} \end{pmatrix} \begin{pmatrix} \nu_L \\ \nu_R^c \end{pmatrix} + h.c. \end{aligned} \quad (3.16)$$

This matrix can be transformed into a diagonal one in order to find the physical neutrino eigenstates and their masses. The calculations will not be performed here but we quote the final conclusions.

Diagonalisation of the mass matrix ends up in two eigenstates ν_1 and ν_2 that are linear combinations of ν_L and ν_R^c . Their mass eigenvalues are:

$$m_{1,2} = \frac{1}{2} \left(m_\nu^R \pm \sqrt{(m_\nu^R)^2 + 4m_D^2} \right) . \quad (3.17)$$

If one makes the assumption that $m_\nu^R \gg m_D$, one can see that the two mass eigenvalues become $m_1 = \frac{m_D^2}{m_\nu^R}$ and $m_2 \approx m_\nu^R$. The second mass eigenvalue becomes equal to the large ν_R mass while the first mass eigenvalue is inversely proportional to the mass of the right-handed neutrino. The Type-I seesaw mechanism therefore offers a natural solution to the small neutrino masses. This mechanism can further be extended to multiple neutrinos.

Without deriving their exact form, we will just pose that the neutrinos in the flavour base are a linear sum of mass eigenstates. The coupling between the flavour neutrinos and the heavy mass eigenstates that follow from the previous assumption of $m_\nu^R \gg m_D$ can be expressed in terms of a coupling matrix V_{lN} . The coupling strength, meaning how large the contribution of the heavy neutrinos to the flavour eigenstates is and thus how likely it is to create HNLs in experiments, turns out to have the proportionality $|V_{lN}|^2 \sim \frac{m_D}{m_\nu^R}$ for a single generation. This means that while the Type-I seesaw offers a natural solution to small neutrino masses, it also comes with large mass values for the HNL states and thus small coupling to the flavour eigenstates, which makes it hard to observe these HNLs in experiments.

To conclude, note what happens when we set m_ν^{RR} to zero. The two eigenstate masses both degenerate into the Dirac mass m_D (with opposite CP phase) and we end up in the scenario of Dirac neutrino masses.

3.2.6 Low-scale models

There exist extensions to the Seesaw model that try to overcome the low mixing and the need for HNLs with touchably high mass. They find a way to place the mass of the HNLs at TeV level or below so that they can be produced in particle colliders and are appropriately dubbed low-scale seesaw models. One example of such models is known as the inverse seesaw model [72].

It starts from the assumption that there is not one type of singlet neutrino but two types per generation, one of which is the familiar right-handed neutrino ν_R alongside a second type X_L . The mass matrix for a base of $(\nu_L \quad \nu_R^c \quad X_L)$ then becomes

$$M = \begin{pmatrix} 0 & m_D & 0 \\ m_D & \mu_R & M_R \\ 0 & M_R & \mu_X \end{pmatrix}. \quad (3.18)$$

The diagonal terms are once again the terms that are responsible for lepton number violation while the others conserve this number. If one makes the assumption that the lepton violating contributions are very small and that $\mu_R, \mu_X \ll m_D \ll M_R$, one finds mass eigenvalues for the light masses after diagonalization of

$$M_L \approx M_D^T M_R^T \mu_X^{-1} M_R M_D. \quad (3.19)$$

We observe here that the smallness of the neutrino mass no longer stems from a large heavy neutrino mass as in the regular type-I model but from the smallness of lepton number violation in the model. If the lepton number violation is small enough, it is very well possible to have heavy neutrino masses of TeV level or below.

3.2.7 State-of-the-art constraints

The playing field for searches or constraints on HNLs is extensive. Depending on the model, HNL masses could range from eV to Grand Unified Theory (GUT) energy scales [73] that are many orders of magnitude larger than we can reach with present-day particle accelerators. While no sign of HNLs was found yet, a plethora of observations provide constraints in the relevant parameter space. Although the HNL is sterile and can not have direct gauge interactions, these observations can rely on their mixing with SM neutrinos. An overview of the experimental and theoretical constraints is compiled below based on references [74, 75].

Limits on HNLs are typically expressed as a function of the mass of the right-handed neutrino and the coupling squared $|V_{LN}|^2$. The latest such results in both the experimental and theoretical landscape are shown in Figure 3.5. In these plots, all filled-in areas are experimentally or

theoretically excluded parameter space while all lines are projections of expected results in the next 10 to 30 years.

The plots start at HNL masses of 100 MeV. Below this mass, there are a few considerations and observations that rule out the existence of HNLs. For masses below 1 eV, **neutrino oscillation data** are the source of exclusion. If HNLs existed at those masses, they would have a larger impact on the neutrino oscillations. Meanwhile in the range of 10 eV to 1 MeV, important constraints on the coupling between electron neutrinos and an HNL come from **neutrinoless double beta decay** [76, 77]. This last type of search is not limited to the quoted range but can provide constraints beyond the TeV mass scale.

The neutrino mass range of 1 MeV to 1 GeV opens up to **peak searches for leptonic decays of mesons**. When charged mesons decay to a charged lepton and a neutrino, processes where the neutrino is the HNL add a monochromatic peak to the energy spectrum of the charged lepton in the decay when observed in the meson rest frame. This idea can be extended up to the use of heavier charged baryons, as done by Belle in 2012 [78], where they looked at $B \rightarrow X\ell N$ the HNL N decaying through $N \rightarrow \ell\pi$. Future B-factories have the potential to also place stringent constraints on the coupling $|V_{\tau N}|^2$, which is currently less-probed than its lighter lepton variants, by producing large amounts of hadronic tau decay events where the tau neutrino mixes with an HNL [79].

A lot of results in the same mass region of up to 1 GeV also stem from **beam dump experiments**. As the name suggests, these experiments dump the beam into material that absorbs all particles from interactions except for neutrinos. Because of their mass, heavy neutrinos will decay and a detector placed at a distance behind the dumping point can measure its decay products. Examples of such beam dump experiments are PS191 [80], NuTeV [81] and CHARM [82–84]. The DUNE experiment [85] (labelled LBNE in Figure 3.5), estimated to be online in the near future, will provide near-detector observations at a length of 30 m of neutrinos produced in charm decays. It has the potential to probe much lower values of the coupling in this mass range than our current observations.

A third type of searches that probe the mass range up to $\mathcal{O}(\text{GeV})$ consists of **searches for lepton number violating decays in mesons** where a meson X decays to $X \rightarrow \ell^\pm N$ and the neutrino of Majorana nature further decays to $N \rightarrow \ell^\pm X_2^\mp$. LHCb has performed such a search in the channel $B^- \rightarrow \pi^+ \mu^- \mu^-$ [86]. The resulting limit is shown in the middle plot of Figure 3.5.

The two previously mentioned types of searches have the disadvantage over the peak searches in the sense that they require the HNL to decay. This means that its decay and its resulting decay products need to be found within the detector volume, which ends up in a suppression of the number of signal events. Experiments like the SHiP [87] will be able to overcome this challenge by highly increasing the flux of incoming hadrons. SHiP will do this by placing a beam-dump facility at the Super Proton Synchrotron (SPS) ring at CERN that will receive a high-intensity proton beam. It will probe decay distances of the order $\mathcal{O}(50\text{m})$. Figure 3.5

shows two projections for this experiment. This is because the production fraction of B_c , which contributes to the HNL production channels, is not measured at the SHiP centre-of-mass energy and is treated as unknown. Two assumptions are made for this cross section: one according to the LHC and one without B_c decay [88].

Two other experiments popping up in Figure 3.5 are called FASER [89, 90] and MATHUSLA200 [91]. Both of them are (or will be) CERN-based experiments. The first of these two, FASER, is an experiment that capitalizes on the loss of particles along the beam direction due to the lack of detectors in the very forward direction at the ATLAS interaction point. By placing a forward detector down the tunnel it can catch forward leptons from light and weakly interacting particles. At the time of writing this section, Run III has just commenced at CERN and the FASER experiment has started taking its first data [92]. While the Run III data will not be enough to be competitive with other sterile neutrino experiments that probe light lepton neutrino coupling, FASER2 [93], expected to start running at the end of 2027, has the potential to improve on the current limits in the mass region up to a few GeV. It will probably be one of the first experiments to follow DELPHI in probing tau neutrino coupling for this mass region. The second experiment, MATHUSLA is an experiment planned for 2025 that places detectors above ground over an interaction point at CERN. It has the potential to provide great improvements to the current limits for light lepton neutrino coupling.

Most of the exclusion results from outside the LHC tend to focus on HNL masses below 1 GeV where the HNLs are long-lived. The lifetime of an HNL is proportional to both its mass and its coupling in the following way: $\tau_N \propto m_N^{-5} V_{\ell N}^{-2}$ [94]. This typically means that for those searches the HNL will travel for a while before decaying. When we start going up to HNL masses of 20 GeV and above, we can start to see that for a large range of mixing parameter values, the mass is high enough for the HNL to have such a short lifetime that it does not get the opportunity to travel and it will decay immediately where it was produced. This is often called *prompt* decay.

When going to these higher mass ranges above $\mathcal{O}(\text{GeV})$, one of the biggest results so far stems from the **DELPHI** collaboration at the Large Electron-Positron Collider (LEP) [95]. Their limits between 1 GeV and 80 GeV masses were obtained by analyzing Z decay events. They performed four different searches: 2 for short-lived HNLs and 2 for long-lived HNLs. From these searches limits on the process $Z \rightarrow \nu N$ were obtained as shown in Figure 3.5. To this day, the results from DELPHI still provide very competitive results in this mass region.

Last but not least there are the direct collider results from experiments like **CMS** and **ATLAS**. These analyses look for signs of on-shell HNLs with masses of the order of the electroweak scale that are directly produced at colliders. The HNL mass range studied by both experiments varies from a few GeV to about 20 TeV. For the lowest masses up to 20 GeV, both ATLAS and CMS have published results using displaced signatures to target HNL signals with non-negligible lifetime. One of these looks at DY-like HNL production in final states with three leptons of which two show displaced signatures [96, 97]. The CMS result places stringent limits on models with exclusive coupling between the HNL and the light-flavour neutrinos

while the ATLAS experiment produced exclusion limits in mixed-coupling scenarios. Both searches have excluded mixing parameter values down to 2×10^{-7} in the mass range of 2 to 20 GeV.

The same channel and final state have also been studied for higher HNL masses without the displaced signatures [98, 99]. The CMS search also targets exclusive coupling to light-flavour neutrinos and is shown in Figure 3.5 while the ATLAS search focuses on coupling between HNL and muon neutrino and put constraints on the mixing parameter down to 10^{-5} for HNL mass values between 4.5 and 50 GeV. Both searches have been performed in a subset of the Run II data. The search in Chapter 7 expands on Ref. [98] by using the full 2016-2018 dataset and by adding results for coupling between tau neutrinos and a promptly decaying HNL. While the displaced search by ATLAS in Ref. [97] includes mixed couplings involving τ neutrinos and CMS has a paper on the way with a displaced search to probe similar mixed couplings in final states with two leptons and a displaced jet [100], so far there were no results available yet for instantaneously decaying HNLs coupling to τ neutrinos.

Other results can be found in the DY production mode with two same-sign leptons and jet final states that violate lepton number conservation [101]. At the high mass range above 700 GeV, CMS and ATLAS have performed searches that target vector boson fusion (VBF) production of the HNL [102, 103], which is the dominant production mode at those masses. The exclusion limits from those searches are the strongest ones available for the muon neutrino coupling scenario in the mass range above 700 GeV with exclusion in the mixing parameter down to 0.1.

There are a few theoretical constraints in Figure 3.5 that have not been addressed yet. One of them is the **seesaw line**. For any mass-coupling point below this line, the mixing of the HNLs with the active – i.e. non-sterile – neutrinos would simply be too small to produce the observed neutrino oscillations in the assumption of one HNL [104, 105]. Similarly, there are constraints that stem from Big Bang Nucleosynthesis (BBN) arguments. If HNLs would couple too weakly to active neutrinos at low enough masses, it would end up in overproduction of Helium-4 in BBN [106]. Finally, there is a band of exclusion obtained by performing a fit to electroweak precision data [107].

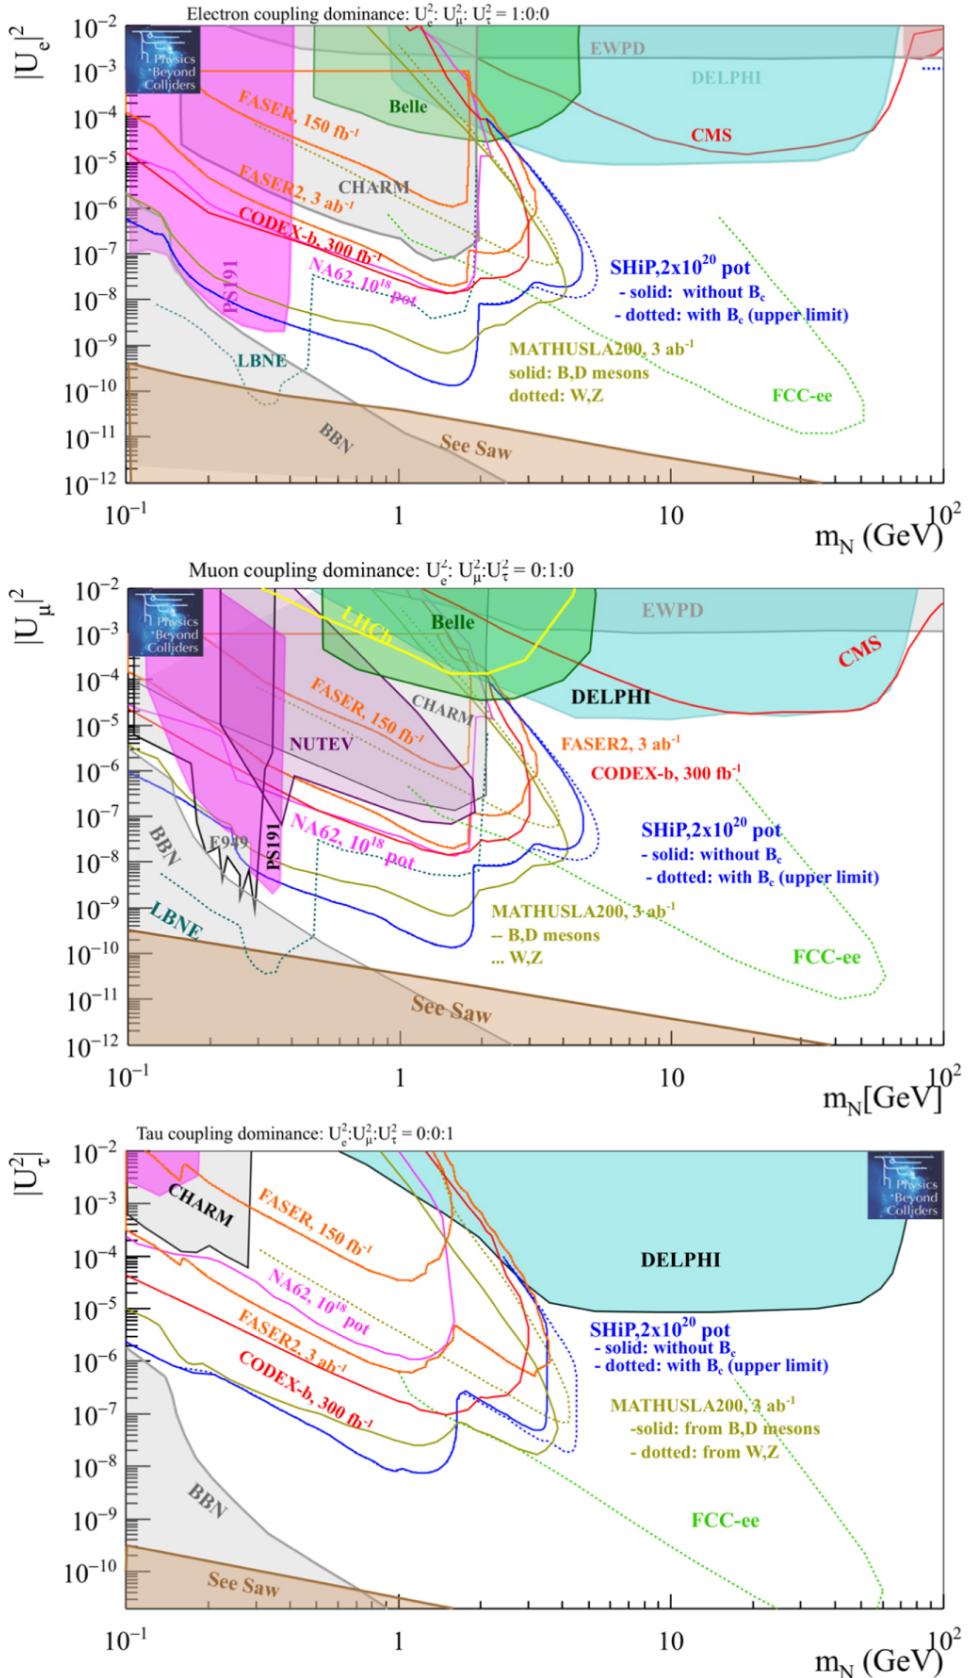


Figure 3.5: Theoretical and experimental constraints on the mass of a single HNL and its coupling to electron neutrinos (top), muon neutrinos (middle) or tau neutrinos (bottom). For each coupling type, only coupling to that generation is assumed to be non-zero [75].

4

The CMS experiment

4.1 The Large Hadron Collider

Even though the SM describes the existence of a range of fundamental particles, only a fraction of them can exist as stable particles in nature. Massive particles will decay to particles with lower mass if allowed by conservation laws. Weak bosons for example will decay to lighter particles on a timescale of the order of 10^{-25} s [23]. Unfortunately this means that these particles can not readily be observed in nature and we need a way of producing them ourselves if we want to perform studies. This is where particle colliders come into play. By accelerating stable particles to speeds near that of light and colliding them, the energy in the collision can be converted into the production of massive or exotic particles at the interaction point. These particles will still decay of course, but their production and decay will occur in an environment under control.

The LHC at CERN is such a particle collider based near Geneva and is designed to accelerate hadrons. Most of the time the accelerated particles will be protons but occasionally heavy ions are injected for special runs. The LHC is the largest particle collider in the world, consisting of a circular accelerator that is 27 km in circumference and was built in the tunnel that housed its predecessor accelerator, LEP, an electron-positron collider. One of the main advantages of accelerating protons instead of electrons and positrons is that protons have a mass that is about 2000 times higher. When particles are forced to follow a circular track, the constant change in acceleration means they will emit *synchrotron radiation*, which will steal some of the original energy and take it away. The probability for emission of such radiation is inversely proportional to the mass of the particle squared, which means that the heavier proton is less prone to losing energy to synchrotron radiation than an electron. This allows to reach much higher centre-of-mass energies in their collisions.

Additionally, protons are composite particles consisting of three valence quarks and a sea of virtual quarks, anti-quarks and gluons. The actual collisions at the LHC will occur not through the protons as a whole but between their constituents (called partons), for example between a quark and a sea anti-quark. This is an advantage for searches for particles of which we do not know the mass yet. Since the momentum of the proton is distributed among the quarks

following *parton distribution functions (PDFs)* [11], the actual momentum of the colliding particles, and thus collision energy, is variable, which makes it perfectly suited to scan a range of potential new particle masses in one dataset. The disadvantage here is that there will be significant hadronic background in the detectors due to QCD processes involving the other partons in the protons. Lepton colliders have a much cleaner signal in that sense. It also means that most of the time the lab frame is not the centre-of-mass frame and there will be residual momentum along the beam direction. This removes conservation of momentum along the beam direction as an effective tool for event reconstruction.

Particles at the LHC are accelerated in different stages before entering the main ring. Protons in particular are generated by stripping electrons from hydrogen atoms using electric fields. They are then launched by a linear accelerator (LINAC) into the circular Booster. Next, they make their way to the Proton Synchrotron (PS) followed by the SPS, both of which are circular accelerators that predate the LHC. In the SPS ring, protons are accelerated up to 450 GeV, at which point they are ready for injection into the LHC. An overview of the path that protons follow during acceleration is given in Figure 4.1.

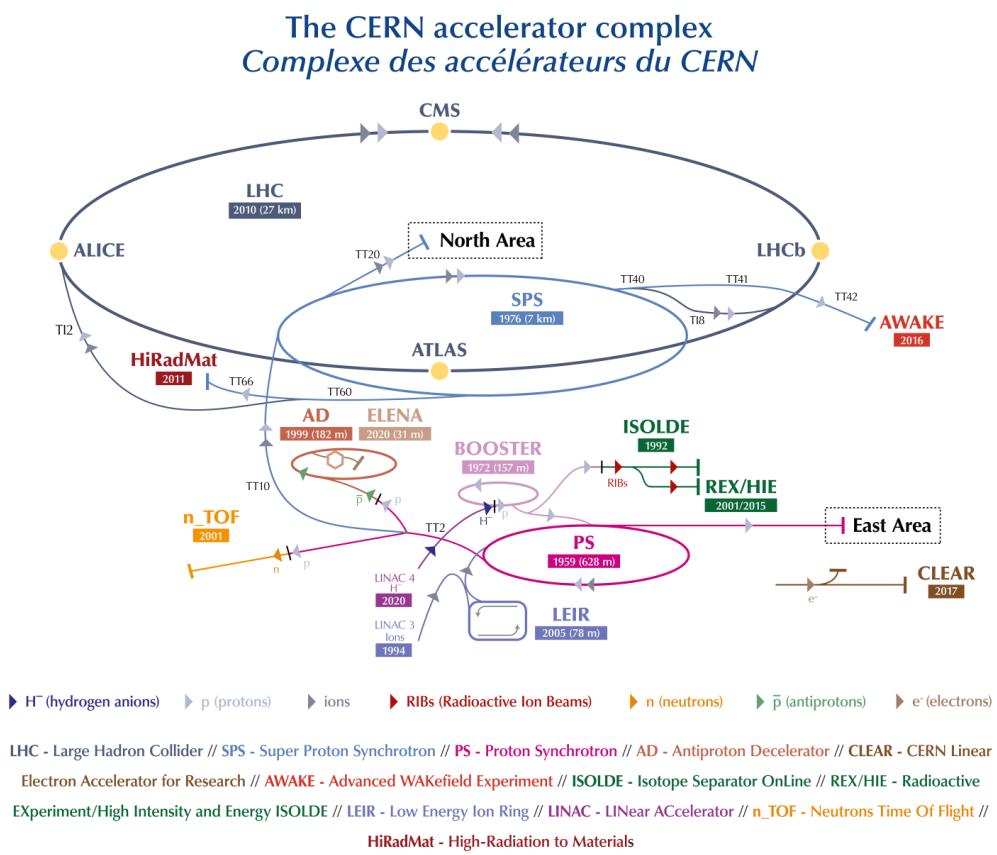


Figure 4.1: Accelerator chain at CERN [108].

The LHC ring consists of two adjacent beam pipes where particle beams are accelerated in opposite directions. Dipole magnets bend the particle beams to follow the circular track while

quadrupole magnets focus the beam. Since accuracy is of the essence, higher-order magnets make small corrections to magnetic field fluctuations. The particle beams are accelerated up to their target energies by special *radiofrequency (RF) cavities* [109], which are metallic chambers containing an electric field with a radiofrequency of around 400Hz that provides the beams with a resonant “push” as they pass through. It takes about 20 minutes for the particle beams to reach their maximum energy after being injected from SPS. At different points on the ring, there are places where the beams can be redirected to collide. Detectors like CMS and ATLAS can be found at such interaction points to observe and study the resulting outbursts of particles.

The output performance of particle colliders can be quantified in terms of their **luminosity** [110], which is a process-independent measure of the rate of collisions at the interaction points that is related to the flux of the beams and is collider-dependent. The expected rate for a specific process can be estimated by combining the luminosity and cross section, the latter being a quantum-mechanical probability for the process to occur.

$$\frac{dN}{dt} = L_{inst}(t)\sigma . \quad (4.1)$$

Here $\frac{dN}{dt}$ is the rate of a process at time t and σ is the cross section of the process. L_{inst} stands for the *instantaneous luminosity*, which is a measure of the luminosity at a specific time t . The target instantaneous luminosity for which the LHC was built is $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. The expected number of events of a certain process in a collider over a period of time can be obtained from the integrated luminosity:

$$N = \int L_{inst}(t)\sigma dt = L\sigma . \quad (4.2)$$

Here L denotes the *integrated luminosity*. The LHC has completed two main periods of data-taking at this point in time. The first run, appropriately labelled Run I, ran from 2010 to 2012 at centre of mass energies of 7 TeV before 2012 and 8 TeV during 2012. It delivered an integrated luminosity of 29.4 fb^{-1} . The conventional unit used here is the barn b of which one unit is equivalent to 10^{-28} cm^2 . The famous H boson discovery was obtained using data from this run. Run II ran from 2015 to 2018 at a centre of mass energy of 13 TeV and delivered a whopping 139.9 fb^{-1} . The highest instantaneous luminosity recorded during this period is $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, twice the target luminosity of the LHC. It has helped us greatly in improving precision measurements, discovering processes with low cross sections predicted by the SM and excluding large parameter space areas of new physics models. The searches described in chapters 6 and 7 make use of data gathered in the years 2016 to 2018 of Run II. On July 5th of 2022, CERN commenced Run III, which is planned to run for 4 years.

One last collider-dependent parameter that should be mentioned is called **pileup**. So-called *in-time pileup* is defined as the number of collisions that occur in a single bunch proton-proton crossing. An increase in luminosity often comes hand in hand with higher pileup as beams are focused better or bunches contain a larger number of protons. Higher pileup comes with the obvious advantage that having more interactions means a higher probability of producing rare

processes. On the other hand common processes like QCD processes will leave traces in the detector alongside the interaction of interest and contaminate its signature. It also means that the detectors will be hit with a colossal amount of radiation and will need to be designed to be radiation-resistant. During the Run II data-taking period, the value for the average pileup was around 30. Because the time for the subdetectors to settle back after observing a signal is non-zero, it is also possible that there will be event contributions from previous bunch crossings to the current beam crossing. This is referred to as *out-of-time pileup*.

4.2 The Compact Muon Solenoid

One of the interaction points located in Cessy, France houses the CMS experiment, a multi-layered detector where every layer is optimized to detect a specific type of particle signature. All the work performed by the author was done so as a part of the CMS collaboration. Therefore, we will take some time here to get into the specifics of the CMS detector.

4.2.1 The CMS coordinate system

Before we get into the different subdetectors that make up the CMS detector, let us first establish the coordinate system that is conventionally used across all CMS analyses. The very centre of this coordinate system is chosen to be the *interaction point* where the two beams in the LHC ring are forced to collide, which is also the centre of the detector. The z -axis - also called the *longitudinal direction* - is chosen as the direction in which the proton beams travel with the positive direction in the counter-clockwise direction of the LHC ring. The x - and y -axis are defined transversely to the beam direction with the y -axis pointing vertically upwards and the x -axis pointing towards the centre of the LHC ring. Together the x and y axes form the *x - y or transversal plane*. The exact definition of these axes can be seen in Figure 4.2.

The direction of a particle measured in the detector is often expressed in terms of two angles with respect to these axes. The first of these is the *azimuthal angle* ϕ , which is defined as the angle a particle makes with respect to the positive x -axis in the transverse plane and is restricted to values between $-\pi$ and π . The second is called the *polar angle* θ , defined as the angle between a particle direction and the positive z -axis. An inconvenient property of the polar angle is that differences in polar angles are not Lorentz invariant under boosts along the z -axis. A preferred variable that conveys this angular information where the difference is Lorentz invariant under such boosts is the rapidity y defined in equation 4.3

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right). \quad (4.3)$$

However, as one can see, this variable is energy dependent. A third variable that depends only on the polar angle, called the *pseudorapidity* η , can be defined as:

$$\eta = -\ln \left(\tan \left(\frac{\theta}{2} \right) \right) = \frac{1}{2} \ln \left(\frac{|\mathbf{p}| + p_z}{|\mathbf{p}| - p_z} \right). \quad (4.4)$$

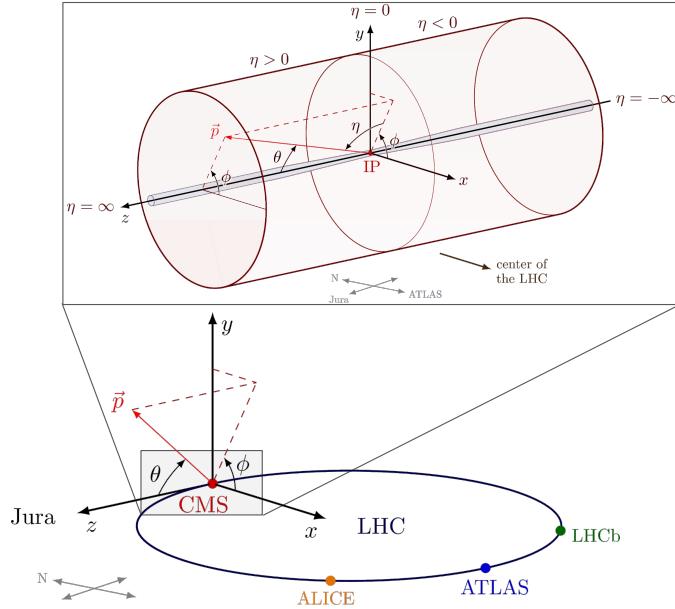


Figure 4.2: The coordinate system at CMS [111]. The bottom part shows the coordinate system of CMS in the context of the LHC ring while the top part in the box shows a zoomed-in version giving a detailed overview of the coordinate system in the context of the detector.

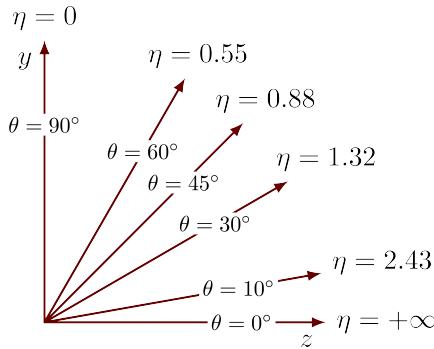


Figure 4.3: The relation between the polar angle and pseudorapidity [112].

In an environment such as the LHC where all particles are highly relativistic and have masses with values much smaller than their momentum, the rapidity and pseudorapidity become the same variable. Therefore, we use pseudorapidity as the main variable to express the direction of particles with respect to the longitudinal direction. To get a better understanding of the relationship between η and θ , the reader is referred to Figure 4.3. A second advantage of the rapidity is that the particle flux is expected to be constant as a function of the rapidity, which is also approximately true for the pseudorapidity.

In the next sections on detector components, you might often see a mention of barrel and endcap regions. The reason why we make such a distinction is that the detector is not infinitely long and is open on the sides. Particles with high pseudorapidity will thus suffer from edge effects in reconstruction. To mitigate this as much as possible, the high-pseudorapidity region is closed off with so-called *endcaps*. While the *barrel* region in the low and medium

pseudorapidity range contains detectors concentric around the beampipe, endcap subdetectors are oriented in the transverse plane so they can target particles with high pseudorapidity. Concrete examples of this can be found in Figures 4.4 to 4.8.

To wrap up this subsection on the coordinate system, I would like to give some attention to a commonly used variable in CMS analyses, namely the *transverse momentum* p_T . When two proton beams meet, their partons interact. The longitudinal momentum of these partons is not fixed, however, but is distributed according to a PDF and the momentum of the two interacting partons is not necessarily the same. Therefore it is likely that the child products of the interaction receive an initial boost along the longitudinal axis of which the size is unknown. Since the protons are accelerated in the longitudinal direction only, the initial partons have only small intrinsic momentum in the transverse plane. The projection of the total momentum vector in the transversal plane is called the transverse momentum vector and its size is called the transverse momentum p_T . Because, unlike the total momentum, the sum of all transverse momenta is conserved if all particles are detected and the transverse momentum is z-boost invariant, it is commonly preferred to use transverse momentum in analyses over the total momentum of particles.

4.2.2 The CMS detector

The CMS detector is made up of layers of subdetectors, each designed to target a different aspect of the signatures left behind by electrons, muons, photons, neutral hadrons and charged hadrons, and a strong solenoid that provides a magnetic field in the subdetectors that bends charged particles in the transverse plane. A cross-sectional overview of the different subdetectors can be found in Figure 4.4, each of which, starting from the inner layers and making our way out, will be the subject of the next few sections.

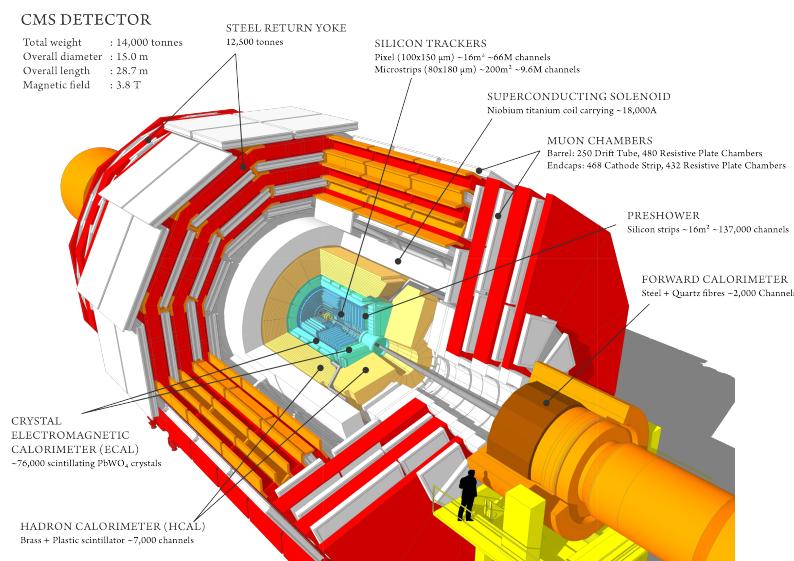


Figure 4.4: The CMS detector [113].

Silicon tracker

Particles produced in collisions will first make their way through the silicon tracker. The main goal of the tracker is to accurately reconstruct the track followed by any charged particle that passes through it. Because of the presence of the strong magnetic field, charged particles will follow a curved path with a curvature defined by the radius:

$$R = \frac{p_T}{qB} , \quad (4.5)$$

where p_T is the transverse momentum of the particle, q is its charge and B is the strength of the magnetic field. This means that we can get a measure of the transverse momentum by accurately measuring the radius of the curvature of the track and the sign of the charge from the direction of this curvature.

The silicon tracker operates on the principle of *ionisation*. Charged particles passing through the submodule create electron-hole excitations which are in turn guided towards electrodes by an electric field. By stacking layers of lightweight material, charged particles leave ionisation hits in each subsequent layer. Linking the positions of the hits in each layer and performing a fit allows for the track of the charged particle to be reconstructed.

The material of choice for the tracker system of the CMS detector is silicon. It is lightweight and thus will allow charged particles to pass through without significant interactions. At the same time, it is radiation resistant, which is an important necessity given the position of the silicon tracker as the first to be greeted by an immense flux of highly energetic particles.

There are two distinct parts to the silicon tracker: the pixel detector and the silicon strips. The pixel detector is the part of the tracker that makes up the first few layers. Three sheets of silicon cubes of size $100 \times 150 \mu\text{m}^2$ make up a barrel layer while the endcaps are made up of two such sheets. To each of these regions, an additional layer was added in the shutdown period between 2016 and 2017 data-taking. The pixel detector allows for a high granularity, which is of utmost importance as the first layers are hit with a lot of decay products that are not yet spread out.

Following the pixel detector are layers of strip modules which operate on the same principle as the pixels but are wider and thus have lower granularity. The barrel area contains 10 such layers distributed over what is called the tracker inner barrel (TIB) and tracker outer barrel (TOB). Contained within the inner barrel are also 6 tracker inner discs (TIDs) while the endcaps are made up of 9 layers of strips. An overview of the silicon tracker can be found in Figure 4.5.

Electromagnetic calorimeter

Whereas the silicon tracker is designed to accurately measure the transverse momentum of charged particles, the calorimeters are designed to measure the energy of particles. The electromagnetic calorimeter (ECAL) specifically targets electrons and photons. The core idea behind the calorimeter is showering of the particles. An electron passing through will emit

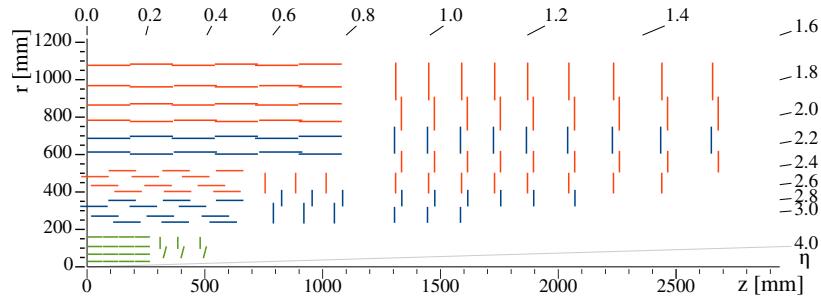


Figure 4.5: A quarter of the silicon tracker in the y-z plane during the 2017-2018 period of data-taking. The pixel detector, shown in green, had one less layer before 2017. The red and blue segments show single-sided and double-sided strip modules respectively. The latter contain two modules that are orientated differently in the plane of the module [114].

bremsstrahlung photons while a photon passing through will interact with the calorimeter material and produce secondary particles. In both cases, the final products will in turn interact with the material and continue the showering cycle. Scintillator material in the calorimeter will get excited by the shower tracks and emit light pulses. By measuring the number of light pulses, one can recreate the energy of the original particle after calibration. Unlike the tracker, calorimeters stop the particles to measure their energy and as such perform a destructive measurement.

The CMS detector uses a layer of lead-tungstate crystals for its ECAL. They have a high density so the electrons and photons interact with them, are optically transparent so the light pulses can travel through, and scintillate when electrons or photons pass through. Each crystal is fitted with a photodiode in the back that measures the produced light and translates it into a signal. The barrel region consists of the pseudorapidity region $|\eta| < 1.479$ while the endcap makes up the region $1.635 < |\eta| < 3.0$. A single crystal in the barrel has a size covering a region $\Delta\eta \times \Delta\phi = 0.0174 \times 0.0174$ (translating to an area of $22 \times 22 \text{ mm}^2$ on the front side of the crystal), allowing for a high granularity. In the endcap region, the size of the crystals is increased to $28.4 \times 28.4 \text{ mm}^2$ on the front side. The endcap area is also fitted with preshower detectors right in front of the crystals with the purpose of making a distinction between high-energy photons and closely-spaced photon pairs from π^0 decay. A geometric overview of the ECAL detector and its components can be found in Figure 4.6.

Hadronic calorimeter

The ECAL is followed by the hadronic calorimeter (HCAL), which serves the purpose of absorbing hadrons and measuring their energy. It is made up of alternating layers of absorber plates and plastic scintillators. The absorber plates will cause hadrons to interact with its nuclei and produce a cascade of particles that keep interacting with the HCAL until they are fully absorbed. When these cascades – also called showers – pass through the scintillator layers, they induce light emission which similarly to the ECAL is measured as a proxy for the hadron energy.

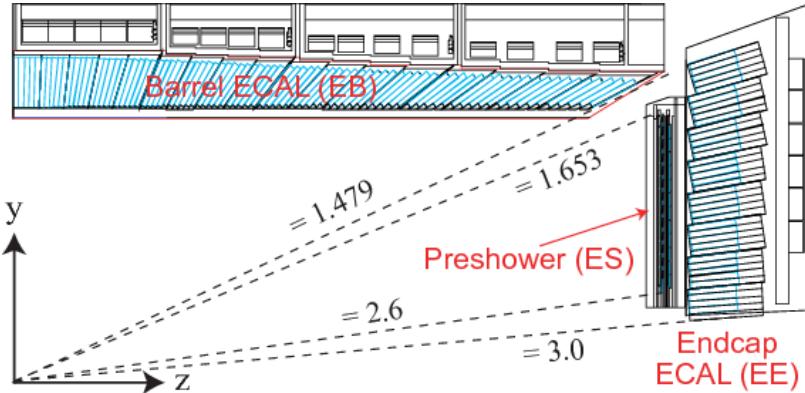


Figure 4.6: A schematic overview of a quarter of the ECAL detector in the r-z plane. [115]

A geometric overview of the HCAL is given in Figure 4.7. The main parts of the HCAL are the HCAL barrel (HB) and the HCAL endcaps (HEs). They cover a region of $|\eta| < 1.3$ and $1.3 < |\eta| < 3.0$ respectively. The absorber material in the HB and HE is brass, chosen due to its short hadronic interaction length of 16.42 cm and because it is ferromagnetically neutral. The granularity of this detector is lower than that of the ECAL with towers of size $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ for $|\eta| < 1.6$ and $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$ for $|\eta| > 1.6$.

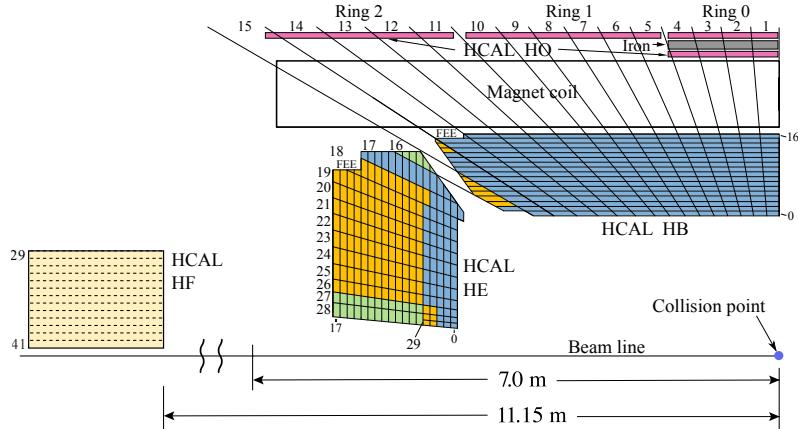


Figure 4.7: A schematic overview of a quarter of the HCAL detector in the r-z plane. [116]

A part of the HCAL is situated outside of the magnet and is called the outer HCAL (HO). Since the CMS detector is such a compact detector, parts of the hadronic showers might make it to the magnet where they can be further absorbed. The HO places scintillators beyond the magnet to catch the tails of these showers. In the pseudorapidity range around 0, where there is less absorber material, these scintillators are preceded by iron plates.

To cover the pseudorapidity range of forward jets close to the beamline, a forward HCAL (HF) is put in place. It covers the range of $3.0 < |\eta| < 5.0$ and makes use of steel as the absorber material. Quartz fibres run through the HF and allow the light to be measured. Here, the electromagnetic component can emit Cherenkov radiation which can in turn be measured.

Solenoid

The magnet is one of the defining features of CMS and is responsible for bending the charged particles so that their transverse momentum can be measured in the trackers. It is a large solenoid built from niobium-titanium coils and was designed to provide an axial, homogeneous magnetic field of 4T inside the solenoid. In reality, the solenoid operated at 3.8 T to increase its lifetime. Outside the solenoid, iron layers are intertwined with the muon trackers, which will be discussed in the next section. We call these layers the return yoke and it guides the magnetic field outside of the solenoid and ensures that a magnetic field strong enough to measure the momentum in the muon detector is maintained.

Muon detectors

It was mentioned that the main way of losing energy and initiating a shower for electrons occurs through bremsstrahlung. While the exact way particles interact with the detector material is non-trivial [23], the rate of energy loss is inversely proportional to the mass of the particle squared. This means that this rate is suppressed for muons by a factor (m_e^2/m_μ^2) compared to electrons. Unlike their lighter counterparts, muons will generally not interact with the calorimeters and make their way further outwards.

In order to detect muons and measure their kinematic properties, dedicated muon trackers are placed outside of the solenoid, interleaved with the return yoke. With the large amount of material in between the beamline and the muon system, the potential for particles that are not muons – with the exception of particles invisible to the detector such as neutrinos – to make it into the muon trackers is very small and therefore muon signatures are quite clean.

Three types of muon trackers were operational during Run II: drift tubes (DTs), cathode strip chambers (CSCs) and resistive plate chambers (RPCs). A geometric overview of all three detectors is shown in Figure 4.8. The first on the list, **DTs**, are tubes 4 cm in width that contain gas and a single wire with a positive charge. When a charged particle such as a muon passes through the tube, it ionizes the gas. The resulting electrons follow the electric field towards the wire and by measuring where along the wire they arrive as well as the distance of the muon from the wire, a position in 2D space can be determined. The distance from the wire can be calculated from the nearly constant drift velocity of the electrons in the gas. DTs cover the barrel area of the CMS detector up to $|\eta| < 1.2$. They are placed in layers where the wires are parallel to each other and 4 such layers are stacked in a zig-zag pattern into so-called superlayers. Four DT chambers are fitted in the barrel, each of which contains 2 or 3 such superlayers. The inner superlayer measures the position along the beamline while the outer layers measure the position in the transverse plane.

Generally, the rate of muons in the barrel will be lower than in the endcaps. To compensate for this, the muon detectors in the endcap area will need to operate faster. The **CSC detectors** that cover the region $0.9 < |\eta| < 2.4$ meet this requirement. They consist of chambers filled with a gas and a range of positive anode wires with orthogonal to them a range of negative cathode strips. When a muon ionizes the gas, the electrons will make their way to one of the

anodes while the resulting ion will drift to the nearest cathode. By measuring the position where the cathode and anode are hit, two-dimensional coordinates for the muon position can be determined. Four detector chambers, each made up of 6 layers of CSCs, make up the CSC section of the muon detector.

The third type of detector in the realm of CMS muon detectors in Run II is the **RPC**. They are important to the trigger system due to their time resolution of $\mathcal{O}(\text{ns})$. The concept behind the RPCs is to have two parallel resistive plates over which a large voltage is applied. These two plates are separated by a layer of gas. Muons passing through the gas will create electrons which will in turn create an avalanche of electrons by interacting with that gas. The avalanche will pass through the plate and will instead be recorded by metallic hit strips attached to the plate. The RPC system is placed both in the barrel and endcap region and covers the region up to $|\eta| < 1.6$.

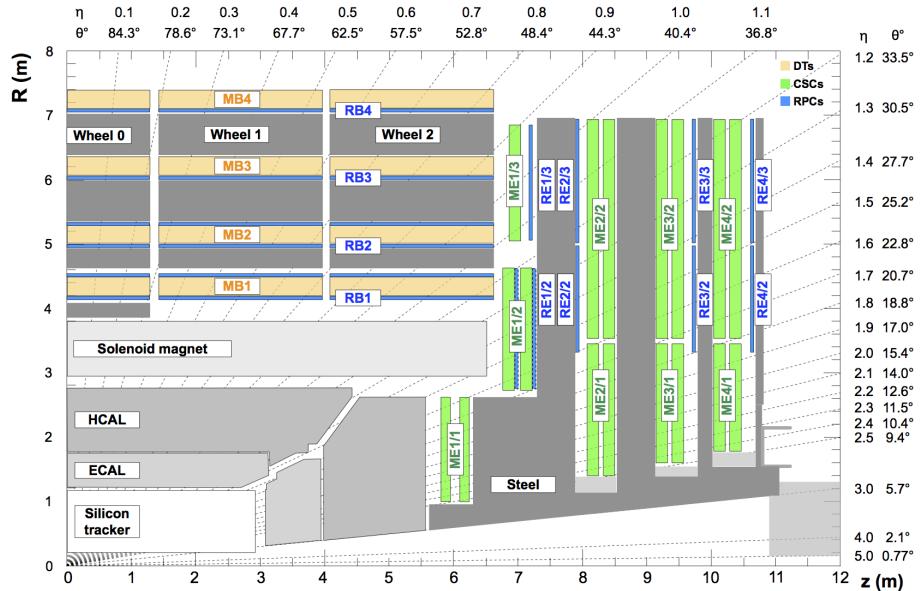


Figure 4.8: A schematic overview of a quarter of the muon detectors in CMS in the y-z plane. DTs are shown in yellow, CSCs in green and RPCs in blue [117].

4.2.3 Trigger system

The LHC has the potential to provide collisions every 25 ns. This leads to an input rate of 40 MHz, which is too high to be able to store all events for offline analysis. Therefore the CMS experiment uses a two-tiered trigger system that selects events of interest to the physics programme for storage. An extensive overview and discussion of the trigger system at CMS will be provided in Chapter 5.

4.3 Object reconstruction in CMS

Particles traversing the detector will interact with the different subdetectors in different ways. Electrons and photons will shower in the ECAL while hadrons will start to shower in the ECAL but will mainly move through to the HCAL and be absorbed there. Electrons and the charged hadrons will leave a track in the silicon tracker as opposed to the photons and neutral hadrons that do not interact with it. Finally, muons will leave tracker hits in both the silicon tracker and the muon system and leave ionisation energy in the calorimeters without showering. The strength of the CMS particle reconstruction comes from combining the information from all subdetectors to construct individual particles in the so-called *Particle Flow (PF) algorithm*. Let us briefly go over the components of particle reconstruction in the following subsections.

4.3.1 Tracking and vertex reconstruction

After charged particles leave hits in subsequent layers of the trackers, the tracking algorithm collects them and uses these hits to construct tracks that describe the particle's original trajectory. This is done iteratively using a combinatorial track finder that builds and fits the tracks [118]. Three main stages make up the procedure. Firstly seeds are built close to the interaction point from a limited number of succeeding hits that show great potential for forming a track. The second step where the actual tracks are formed starts from these seeds to build tracks using a *Kalman Filter (KF)* [119]. The initial seed track is extrapolated and a hit in the following layer that is consistent with the track is added if available. After adding a hit, the track is refitted and the new track parameters are determined. This procedure continues until the final tracker layer is reached. At the end of the procedure, a final fit is performed to the track from which the track parameters are determined. Tracks failing a goodness-of-fit test are rejected.

To improve the efficiency of the algorithm an iterative strategy was adopted. High-energy particles that originate close to the interaction point are the easiest to handle with the fitting procedure as they form well-defined tracks that have low curvature. Once these types of tracks have been reconstructed, the algorithm moves on to the next iteration of tracking reconstruction. When a track has been reconstructed, the hits associated to it are masked, meaning they will no longer be regarded for the reconstruction of other tracks. This allows harder tracks such as displaced tracks or low-energy tracks to be reconstructed in a less crowded environment.

Once all tracks have been defined, one can move on to vertex reconstruction. A *vertex* is defined as a point of origin for events where all tracks of a certain event come together. For this step, all high-quality tracks that originate from the beam spot – defined as the luminous region of the LHC where the beams interact – are selected and the z-coordinates of the track origins are collected. An *adaptive vertex fitter* [120] first determines the positions of potential vertices and assigns tracks to them. At the same time, a weight is calculated for each track, which quantifies the probability that the tracks actually originate from the vertex. If the weighted sum of all tracks in a certain vertex does not beat a certain threshold, the vertex is rejected.

Once all vertices have been reconstructed a single vertex is chosen to be the *primary vertex*, being the vertex with the highest capacity to be an interesting event in terms of the physics programme of CMS. In reality, this choice is made by selecting the vertex corresponding to the hardest scattering in the event as described in Section 9.4.1 of Ref. [121].

4.3.2 Calorimeter clustering

Most particles will be absorbed in either of the calorimeters. Constructing the energy deposits from single particle candidates allows us to determine their energy and direction, separate different contributions in jets and improve the electron resolution by measuring the contributions due to bremsstrahlung photons [118]. The calorimeter clustering algorithm will start by selecting seeds. In the context of the calorimeters, these seeds consist of cells where the energy deposit passes a certain threshold and is higher than that in the neighbouring cells. Secondly, so-called *topological clusters* are constructed by adding cells that share one corner with the cluster and where the energy deposit outweighs the noise by a factor two. Finally, a Gaussian mixture model algorithm fits different clusters – different particles – to this topological map. The calorimeter clustering is performed separately for the ECAL barrel and endcaps, the HCAL barrel and endcaps, the HF and the two preshower layers.

4.3.3 Particle Flow

With all tracks, vertices and calorimeter clusters in hand, the PF algorithm will start to link matching elements into PF blocks. From these blocks, individual particles are reconstructed, after which the elements that were used are removed from the block. By exploiting the knowledge of the intrinsic energy and momentum resolution of each subdetector, the PF algorithm can perform this combination in the most accurate way. To get an in-depth idea of the ins and outs of PF, the reader is referred to external sources [122]. The algorithm outputs five different particle types, namely electrons, photons, muons, charged hadrons and neutral hadrons. They form the basic building blocks with which more complex objects such as hadronized quark or gluon jets and τ leptons are reconstructed. It should be noted that the tracking for electrons and muons contains a few extra steps as opposed to the basic tracks discussed before.

4.3.4 Muon reconstruction

Muons in CMS have the unique advantage that they are the only charged particles to pass through both the tracker and muon detectors. The silicon tracker is a highly optimized tracker that can provide us with high-precision momentum and direction measurements. The muon detectors are protected from most particles due to absorption in the calorimeters and thus have a high purity and identification efficiency for measuring muons. By combining the information of both subsystems, one can get a highly pure and efficient muon reconstruction. Three types of muon track candidates are defined.

The first of these are the **standalone muons**. They are muon track objects built purely from hits in the three muon detector types. The seeds for these types of track candidates are

generated from hits in the DT or CSC detectors. The full standalone muon tracks are then built from a fit of these seeds to hits in the DTs, CSCs and RPCs.

The second type consists of **global muons**. By starting from standalone muon tracks and extrapolating to the silicon tracking system, one can try to find a matching track. When there is such a match between a standalone muon track and a track in the silicon tracker, a fit is performed to the hits from both subdetectors. This results in a higher momentum resolution than can be obtained with the separate subdetectors for muons with a p_T over 200 GeV [117]. When a muon penetrates multiple planes of the muon system, they tend to be reconstructed as global muons.

When a muon has a low p_T of the order of 10 GeV, it might not be properly reconstructed as a global muon due to scattering in the return yoke. In this case, the third type of muon tracks, the **tracker muons**, tend to recover them. Whereas global muons start from fully reconstructed standalone muons and work their way inwards, tracker muons start from silicon tracks and make their extrapolation outwards to try and find a matching DT and/or CSC segment. The tracker muon is then defined as the corresponding track from the silicon tracker.

When a global muon and a tracker muon share the same track in the silicon tracker, the two candidates are merged into a single candidate. The three types of muons discussed here are then all given as input to the PF algorithm, which has different conditions for each type of muon.

4.3.5 Electron reconstruction

Electrons will leave a track when they traverse the tracker and will then shower in the ECAL where they are absorbed. Two strategies for reconstructing electrons are applied and combined by the PF algorithm to produce the final collection of electron candidates. The first starts from energetic ECAL clusters and is appropriately called the **ECAL-based approach**. From the ECAL cluster position and shape, one deduces the location of origin in the innermost tracker layers. Most electrons will emit bremsstrahlung while traversing the tracker, which will propagate to the ECAL and produce their own clusters. The ECAL-based approach tries to account for that by forming a supercluster from all clusters within a certain window around the electron direction. This window is small in the η direction but is made to be wider in the ϕ direction to account for bending in the magnetic field.

The ECAL-based approach works well for electrons that emit no bremsstrahlung or emit only a small amount of energy in bremsstrahlung but less so for electrons that lose large fractions of their energy to bremsstrahlung. To account for the electrons that are missed by the ECAL-based approach, a **tracker-based approach** is introduced. The energy loss of electrons radiating bremsstrahlung photons is highly non-Gaussian and a new fitting algorithm is used to produce electron seeds. It uses all tracks from the iterative tracking phase and makes a preselection on χ^2 of the track and the number of missing hits. Electrons that emit low- p_T photons tend to have high χ^2 while the emission of high- p_T photons causes sudden

changes in electron p_T and thus track direction that causes tracks to be reconstructed with a high number of missing hits. The selected tracks are then refitted using a gaussian-sum filter (GSF) which does not rely on a Gaussian energy loss but as the name implies makes use of sums of Gaussians [123]. It can take into account the sudden direction changes due to photon emission. The ECAL-based and tracker-based seeds are combined and a final fit is performed to get the final candidate.

4.3.6 Jet clustering

Because of colour confinement, coloured particles can not exist in isolation and will hadronize. Quarks and gluons produced in events in the CMS detector are no exception and will produce a shower of hadronization particles. We call this collection of particles a jet. In order to accurately reconstruct information regarding quarks or gluons, one must be able to reconstruct the jets resulting from their hadronization.

In CMS this is done using what is known as the anti- k_T algorithm [124]. The clustering is performed using a custom distance measure d_{ij} between any two PF candidates i and j in the event. It is defined as follows:

$$d_{ij} = \min(p_{T,i}^{-2}, p_{T,j}^{-2}) \frac{\Delta_{ij}^2}{R^2}, \quad (4.6)$$

with $p_{T,i}^{-2}$ the squared inverse transverse momentum of candidate i and $\Delta_{ij}^2 = (y_j - y_i)^2 + (\phi_j - \phi_i)^2$ where y is the pseudorapidity and ϕ is the azimuth. A similar parameter is defined to describe the distance between a particle and the beam:

$$d_{iB} = p_{T,i}^{-2}. \quad (4.7)$$

The values of the distances d_{ij} and d_{iB} are calculated for all particle or jet candidate combinations. If the smallest value is the distance between two candidates, the two are combined into a single candidate and all values are recalculated. If the smallest value is the distance between a candidate and the beam, that candidate is labelled as a jet¹ and not considered for further clustering. This algorithm continues until there are no more candidates left. At CMS, a value of 0.4 is used for R .

The strength of this algorithm comes from the inverse dependence of the candidate p_T . It means that the distance between soft particles² will be much larger than that between a soft and a hard particle. The anti- k_T algorithm will therefore cluster the soft particles around a hard particle into the hard particle before clustering the two soft particles themselves. This will make the algorithm collinear and infrared safe and allows great robustness against soft radiation.

¹The name jet here is used for the purpose of the algorithm but could exist of singular particles such as electrons.

²Soft particles refer to particles with low transverse momentum. At CMS, typical energy scales for low-momentum particles are found around transverse momentum values of a few GeV to a few 10 GeV, depending on the context.

Particles coming from pileup events can end up being part of the jet clusters by accident. In order to remove the pileup contributions, a few countermeasures are taken. Firstly, charged particles that are used in the reconstruction of a pileup vertex are removed from the jet. Remaining pileup contributions such as neutral hadrons are then estimated by studying the p_T distributions of jets in simulation with and without pileup. These contributions are then also subtracted.

Reconstructed jets need to be calibrated so that the measured jet energy corresponds to the true jet energy as closely as possible. This is done using so-called *jet energy corrections (JEC)* that first correct the jet response both in observed and simulated events based on measurements performed in simulation and then correct the jets in data for residual differences with jets in simulation. Another set of scale factors applies smearing to jets in simulation to account for the worse *jet energy resolution (JER)* in data. We will not go into the details of these two calibrations, the interested reader is instead referred to external references [125, 126].

4.3.7 Heavy Flavor identification

Processes that include the production of top quarks are an essential part of the CMS physics programme. Because of their high mass, they will decay before forming a hadron. In more than 99% of the cases, this decay will result in the creation of a b-quark, which will then hadronize and shower. In order to have a good grip on top quark production, but also on channels of importance to studies in the BEH sector, such as $H \rightarrow b\bar{b}$, it is vital that we can identify jets resulting from b-quarks. The first will be especially important in the analyses in chapters 6 and 7 in order to reduce the background contributions due to processes including one or more top quarks. This process of identifying b-jets is called *b-tagging* [127].

B-tagging algorithms in CMS make use of the properties of b-hadrons produced in the hadronisation of b-quarks. Firstly the lifetime of b-hadrons is around 1.5 ps, which means that depending on their momentum, they first travel up to 1 cm before decaying. This long lifetime means that we observe a secondary vertex in the jet from the decay of a b-hadron as shown in Figure 4.9. Additionally, b-quarks have a larger mass than up or down quarks and a harder hadronisation which will result in harder decay products. In around 20% of the cases, a charged lepton will be present in the b-hadron decay.

All of the abovementioned properties are used in machine learning algorithms that produce an output score assigning a probability that a jet is a b-jet and not a jet from a different source. The two main algorithms that are also of important to the analyses in later chapters are called DeepCSV [127] and DeepFlavor [128].

4.3.8 Tau Leptons

Because τ leptons have a relatively high mass ($m_\tau = 1.78$ GeV), they are inherently unstable particles with a mean lifetime of 2.9×10^{-13} s, which is smaller than the timescale of event propagation in the CMS detector [23]. Therefore τ leptons in the detector can not be observed

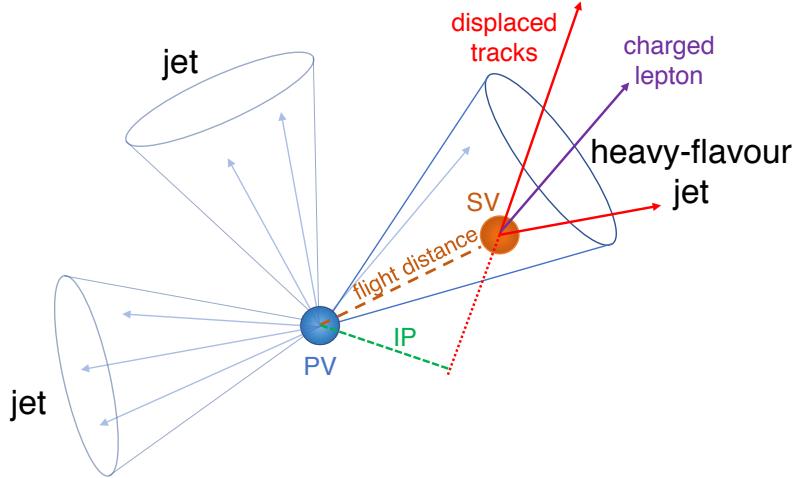


Figure 4.9: Illustration of a b-jet with its primary vertex (PV) and secondary vertex (SV). The figure also illustrates the impact parameter (IP) variable, defined as the distance between the PV and the closest approach of the extrapolation of a particle’s trajectory. [127].

directly but can only be reconstructed from their decay products. Table 4.1 lists the main decay modes for τ leptons along with their according decay products and branching fraction. One of the first things to notice in this table is that there are two types of decay modes: leptonic decays and hadronic decays. The first end up in electrons or muons as well as the needed neutrinos to conserve lepton number, causing the decay products to be nearly indistinguishable from light leptons from other sources. Hadronic decay modes on the other hand look like jets with specialised features such as low multiplicity, lending themselves as much easier candidates for τ tagging. As such we only use hadronic decay modes for τ reconstruction, labelled as τ_h . From now on we will use the term τ lepton in the context of CMS only in reference to a hadronically decayed τ unless the decay mode is explicitly specified.

Hadronically decayed τ ’s are reconstructed at CMS by using the *Hadrons-Plus-Strips (HPS) algorithm* [129–131]. It starts from reconstructed jets as described in the previous section and defines a cone-shaped region in η and ϕ with size $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} = 0.5$ around the jet axis. A few different steps address the different components of hadronic decay as seen in Table 4.1, namely τ neutrinos, charged hadrons, and neutral pions π^0 . The first of these will go undetected as they do not interact with the detector. The charged hadron component will at this stage have been reconstructed by the PF algorithm and are contained in the jet as is. That leaves us with the π^0 component. In 99% of the cases, a neutral pion will decay into two photons, which in turn can convert to e^+e^- pairs. The “Strips” part in the name of the algorithm now has the job to reconstruct these photon pairs by reconstructing so-called strips of electromagnetic activity in the jet, starting from the most energetic electromagnetic particle as the centre of the strip and aggregating all other nearby electromagnetic particles that pass a distance condition into it. It then continues to build a new strip from leftover energetic electromagnetic particles. All particles that are in the end not contained in these strips will no longer be considered part of the τ jet and will instead be used for isolation and identification

Table 4.1: Main τ decay modes, their meson resonance and branching fraction \mathcal{B} [23]

| | τ Decay Mode | Meson resonance | $\mathcal{B}(\%)$ |
|----------------------------------|---|-----------------|-------------------|
| Leptonic Decays | $\tau^- \rightarrow e^- + \bar{\nu}_e + \nu_\tau$ | | 17.82 |
| | $\tau^- \rightarrow \mu^- + \bar{\nu}_\mu + \nu_\tau$ | | 17.39 |
| Hadronic decays (One prong) | $\tau^- \rightarrow h^- + \nu_\tau$ | | 11.51 |
| | $\tau^- \rightarrow h^- + \pi^0 + \nu_\tau$ | $\rho(770)$ | 25.93 |
| | $\tau^- \rightarrow h^- + 2\pi^0 + \nu_\tau$ | $a_1(1260)$ | 10.81 |
| Hadronic decays (Three prong) | $\tau^- \rightarrow 2h^- + h^+ + \nu_\tau$ | $a_1(1260)$ | 9.80 |
| | $\tau^- \rightarrow 2h^- + h^+ + \pi^0 + \nu_\tau$ | | 4.76 |
| Other hadronic decay modes | | | 1.8 |

purposes.

In order to mitigate overestimation effects in the isolation variables due to particles moving outside of a strip to which they should belong when using a fixed strip size, the processing in Run II software made use of a dynamic strip reconstruction algorithm. First, the electron or photon with the highest p_T is chosen to seed the strip. The next step then looks for the highest electron or photon that is within a range:

$$\begin{aligned} \Delta\eta &= f(p_T^\gamma) + f(p_T^{\text{strip}}) \\ \Delta\phi &= g(p_T^\gamma) + g(p_T^{\text{strip}}), \end{aligned} \quad (4.8)$$

where p_T^γ is the photon/electron p_T and p_T^{strip} is the p_T of the strip. The functions f and g are optimized in such a way that 95% of all electromagnetic τ_h decay products are contained within a strip. Every time a new particle gets added to the strip, the position of the strip is recomputed as a p_T -weighted average of its particle content. If there is no further electron or photon within the window determined by eq. 4.8, the current strip reconstruction reaches its end and a new strip clustering starts if a new electron or photon is available.

After all components have been gathered, the HPS algorithm proceeds to the reconstruction of τ_h candidates. This is done based on the main decay modes as seen in Table 4.1. One-prong decay modes are targeted with τ_h candidates containing one charged hadron and up to two π^0 's while three-prong decay modes are targeted with three charged hadrons and up to one π^0 . The absolute value of the total charge of the candidate must be equal to one for these decay modes. Lastly, in more recent years there is also support for decay modes with two charged hadrons and up to one π^0 . These target boosted three-prong decays where one of the charged hadrons is lost. It is important to note that while the recovered three-prong decays with these two prong decay modes are not negligible, it suffers from a reduced charge assignment efficiency. For that reason, they are recommended for use only in analyses that are not limited by low

background event yields. Because none of the analyses described in this thesis satisfy that requirement, these decay modes are vetoed in the analyses described in chapters 6 and 7.

As a last step, a single τ_h candidate is chosen from all the candidates reconstructed in the previous step. A few conditions are taken into account to make this decision. One of these is the mass consideration: the candidate must have a mass that is compatible with either the according resonance of their decay or in the case of one-prong decays the mass of charged pions. Secondly, all charged hadron and strip components must be inside of a predefined signal cone. If multiple candidates pass all conditions, the candidate with the highest p_T is chosen as the reconstructed τ_h .

4.3.9 Missing energy

Neutrinos and other non-interacting particles will generally traverse the entire detector without interacting with it and thus without leaving a signal. This makes it impossible to directly observe and measure their momentum. We do, however, have an indirect handle on their transverse momentum. Momentum conservation dictates that the total transverse momentum vector of all particles to come out of an event is zero but neutrinos will carry away a part of this momentum. Therefore, by inversely summing the transverse momentum vectors of all observed particles in an event, we have a measure of the p_T of all missing particles:

$$\vec{p}_T^{\text{miss}} = \sum_i \vec{p}_{T,i} . \quad (4.9)$$

This inverse vector sum is referred to as the *missing transverse momentum vector* and its magnitude is denoted as p_T^{miss} . Whenever we talk about the missing transverse momentum (MET), we will be referring to the magnitude.

4.4 Event simulation

All components so far have shown us how the CMS detector can record high-energy collision processes for us to analyse. In order to do any meaningful analysis of the observed data, however, we will need a handle on what we expect to see, such as signal and background predictions. In the bulk of the cases, the way to go about this is by simulating the processes that would take place in proton-proton collisions and how they interact with the CMS detector. Typically one does not go about simulating general proton-proton collision events including all possible processes and final states as we would have in reality. If we did, the main part of the resulting processes would be inelastic scattering events that will be thrown away in processing anyway. In order to perform any useful predictions we would need to simulate a colossal number of events. Additionally, it would take a lot of processing time to sift through all of them in post-processing. What is done instead is to simulate a specific process, such as the hard scattering processes of interest. By simulating the most important contributions to an analysis and putting the samples together, one can achieve the same end result as a general simulation.

At CMS these simulations are performed using Monte Carlo (MC) techniques [132, 133]. It should be noted that the full simulation chain and the theory behind it is a highly complex subject. The interested reader wanting to learn more about it is referred to external sources such as Ref. [134]. In order to provide some context for some of the more technical parts of the analyses discussed later on, I will spend the rest of this section on giving a brief introduction to the three main steps of generating simulations. These consist of matrix element (ME) level calculations to simulate the core of the hard scattering process, the further decay and hadronization of these initial particles and lastly their interaction with the detector.

4.4.1 Matrix element-level calculations

The first steps of event simulation start with simulating the core process of interest. The ME in the name actually refers to the scattering matrix element in QFT that links the initial and final states and that is used to calculate experimental observables such as the cross section and decay width of a process resulting from parton interactions [135]. It is possible to express this scattering matrix in terms of the Feynman amplitude which in turn can be expressed as a sum of Feynman diagrams. In fact, this is exactly what was depicted in Figure 2.4. I realize that some of the terms used in the last few sentences come out of nowhere for the inexperienced reader. The main takeaway is that the raw, hard scattering processes between partons involve high momentum transfers. The same idea behind asymptotic freedom as described in Section 2.2.1 now also states that at the considered momentum transfers in these processes, the strong coupling constant α_S is relatively small (at least smaller than one). The Feynman amplitude could then be calculated in a perturbation theory in terms of the strong coupling constant with each order adding an extra strong vertex and thus an extra factor α_S , therefore making their contributions to the sum smaller with each order.

It is important to remember that these ME calculations mentioned so far describe the cross section of an interaction between partons. At the LHC, however, we do not collide isolated partons but protons. The partons inside a proton no longer reside in the same high-energy regime and their behaviour can not be calculated perturbatively. In practice, the proton-proton cross sections for processes at the LHC are calculated using the factorisation scheme that states the high-energy and low-energy regimes can be treated separately [136, 137]. In simplified form, the proton-proton cross section can be written as:

$$\sigma = \sum_{i,j} \int_0^1 dx_i \int_0^1 dx_j f_i(x_i, \mu_F) f_j(x_j, \mu_F) \hat{\sigma}_{ij}(\mu_F, \mu_R) , \quad (4.10)$$

where the sum runs over parton types, x_i represents the fraction of total hadron momentum carried by the parton, f_i is the PDF which describes the distribution of x_i in a hadron for a certain parton and $\hat{\sigma}_{ij}$ is the hard-scattering cross section between the two partons. The last contains the asymptotically free part of the interaction and can be calculated using perturbation theory as described before. The PDFs however contain the soft part of the confined partons and can not be calculated perturbatively. Instead, it has to be parametrized from measurements. One last thing to point out is that in order to remove singularities from the calculation, f_i and

$\hat{\sigma}_{ij}$ are made to be a function of what are called a *factorisation scale* μ_F and a *renormalisation scale* μ_R . The choice of these scales is somewhat arbitrary but the PDFs and hard-scattering cross sections will vary with varying scales. They will return as some of the most prominent sources of uncertainty in the analyses that follow. Also, the exact choice of the PDFs will be an important source of uncertainty.

In short, the ME level event simulation is two-fold: to compute the differential scattering MEs and to integrate them over the phase space of all involved particles. This way the cross section is calculated and it can be used for MC event generation.

4.4.2 Decays and hadronisation

After the ME level simulation comes the further decay of unstable particles and hadronisation of parton showers. This latter occurs when there is a coloured particle in the final state of the hard process or one is radiated off of one of the final state particles. In either case this coloured particle will end up in a cascade due to colour confinement. At every step, the energy of the cascade components is lowered until one reaches a region where perturbation theory breaks down. In the simulations, this showering is simulated in a probabilistic way to produce new particles until the nonperturbative stage is reached. The coloured particles then hadronize into colourless particles, a process which is simulated from phenomenological models.

Additional initial or final state radiation is added in this step as well. To make sure there is no double counting with radiation from the ME level calculations, matching schemes remove radiation that was added in the previous step. Aside from the hard-scattering process, other soft interactions between partons may occur and the remnants of the proton – which are now coloured – will hadronize as well. The collection of both effects is called the underlying event. As they are also in the nonperturbative region, their modelling is also performed phenomenologically.

4.4.3 Detector simulation

The core process and its decays are now fully simulated. What we are missing at this point is its behaviour in the CMS detector. In all samples used throughout this thesis, a program called GEANT4 [138] is used in their production to fully model the CMS detector. It creates a geometric model of the detector and takes into account the different materials and their position and density, the interaction of particles with said material, the detector response and readout, and even imperfections in the detector. The result is a signal readout in the same format as we would get in measured data. This can be used in the exact same way as the data so we can make a fair comparison. It should be noted that detector simulation can never be 100% accurate. This is the reason why we need to perform detailed data-to-simulation comparison and perform all the corrections mentioned in the next chapter to the simulated samples.

4.5 Statistical tools

In new physics searches such as the ones described in Chapters 6 and 7, it is paramount to have procedures to quantify potential excesses of events over the SM prediction or place exclusion limits on the cross section of a model. Analyses at the LHC typically share the same statistical methodology, which allows for consistent comparison of results [139]. This section will provide a small introduction to this specific methodology for use in later chapters.

Take an analysis that produces output histograms containing N bins where the predicted signal yield, the predicted background yield and the observed yield in bin i are denoted by s_i , b_i and n_i respectively. Considering that we expect each bin to follow a Poisson distribution, we can write down the likelihood to measure the observed yields given the predicted signal and background yields as follows:

$$\mathcal{L}(r) = \prod_i^N \frac{(rs_i + b_i)^{n_i}}{n_i!} e^{-rs_i - b_i}. \quad (4.11)$$

This is called a *binned likelihood function*, in which we have introduced a dependency on the signal strength modifier r . This last variable modifies the cross section of the signal. A value of zero means that no signal is expected while a value of one means one predicts signal yields corresponding to the input cross section. The higher the value of the binned likelihood, the better the agreement between data and prediction. Realistically, an analysis will come with systematic uncertainties, which are implemented as so-called nuisance parameters θ in the binned likelihood function as additional factors in equation 4.11. These nuisance parameters are allowed to vary around their nominal values according to probability distribution functions defined by the corresponding uncertainty in order to maximise the binned likelihood function $\mathcal{L}(r, \theta)$. This essentially performs a fit of the data to the prediction taking into account all uncertainties on the yields.

This binned likelihood function can be used to construct a *test statistic* q_r , which in turn can be used to gauge the compatibility between data and a specific hypothesis. It is defined as follows:

$$q_r = -2 \ln \frac{\mathcal{L}(r, \hat{\theta}_r)}{\mathcal{L}(\hat{r}, \hat{\theta})}, \quad (4.12)$$

where $\hat{\theta}_r$ in the numerator is the set of nuisance parameter values that maximises the likelihood for a specific value of r , and \hat{r} and $\hat{\theta}$ in the denominator are the combination of signal strength and nuisance parameter values that maximise the likelihood globally. The higher the value of this test statistic, the worse the compatibility between the data and the hypothesis for r . One can now determine the probability distribution function of q_r under the signal+background hypothesis and the background-only hypothesis ($r = 0$). This can be done by generating pseudo-data using the best-fit nuisances $\hat{\theta}_r$ and $\hat{\theta}_0$ for the corresponding hypothesis. In later chapters, however, you might see mention of the asymptotic approximation, which means these probability distribution functions are approximated using functional forms [140].

In a last step, p-values are estimated from these probability distributions functions as the probability for q_r to be larger than the observed q_r^{obs} :

$$\begin{aligned} p_r &= P(q_r \geq q_r^{obs} | \text{signal + background}) \\ 1 - p_b &= P(q_r \geq q_r^{obs} | \text{background-only}) \end{aligned} \quad (4.13)$$

A confidence level can then be estimated as the ratio of these two p-values

$$\text{CL}_s = \frac{p_r}{1 - p_b} \quad (4.14)$$

A signal hypothesis is then said to be excluded with a certain confidence level (CL) X if $\text{CL}_s \leq X$. In order to find the 95% upper limit on the signal strength, r is varied until $\text{CL}_s = 0.05$. Similarly, we quantify an excess not by estimating the probability distribution function of q_r , but by estimating the probability distribution function of q_0 under the background-only hypothesis and calculating a p-value based on it.

5

Triggers at CMS

5.1 Personal Foreword

As part of my membership in CMS, I, like all members of CMS, have performed service work within the experiment. Most of this was concentrated around work regarding triggers developed and maintained by the Tau particle object group (POG). In 2018, I started out by optimising the strategy for fitting trigger scale factors, which was followed by measuring the scale factors for all Run II triggers and studying the effect of correlation between the two trigger candidates in double τ triggers among other smaller tasks over the course of the next two years. At the beginning of 2020, I was approached by the Tau POG conveners with the question, “Would you be interested in taking over from the previous L3 τ trigger convener after their mandate ends?”. As you might have guessed from the existence of this chapter, I humbly accepted the offer and started my two-year term as L3 convener on September 1st 2020.

To briefly sketch the role of an L3 trigger convener, my duties came down to attending the general trigger studies group (TSG) meeting and reporting anything relevant to the Tau POG group, supervising the validation of τ triggers in new CMS Software (CMSSW) releases as well as data quality monitoring (DQM) during data-taking, leading the developments of new algorithms and paths for the Run III period, providing trigger corrections and keeping up with any other matter τ trigger related. It should be noted that these responsibilities refer to the second part of the two-tiered trigger system, namely the high-level trigger (HLT). The first step, called the level-1 (L1) step, is handled by a different group. At any given time this convenership is shared between two people who usually bring a different manner of expertise to the table. As my most direct contributions were to the management and development of Run III τ triggers, the focus of this chapter will be on that aspect of my work. My mandate ended at the end of August 2022, shortly after Run III data-taking commenced. A lot of components in the trigger definitions are subject to updates over time or need to be quickly adapted once the run starts and unexpected or unwanted behaviour is observed. Therefore we will restrict ourselves to the development up to the beginning of Run III. Lastly, I would like to mention that this chapter is not meant to give a detailed, technical overview of the Run III

trigger menu development. Instead, it will be kept short and provide a conceptual overview of the type of work that was performed.

5.2 Introduction

As mentioned in Chapter 4, the LHC has the potential to create collisions every 25 ns. If every collision event was to be reconstructed and stored, with about 1 MB of data per event on average, this would lead to a data influx of around 40 TB/s. It is clear that with a limited bandwidth of around 1000 full events per second, this amount of data is over the capacity of what can be promptly reconstructed. Fortunately – or unfortunately, depending on your point of view – a large fraction of the collision events at the LHC are not interesting to store when viewed in line with the physics case of CMS. Processes that are relevant for new physics analyses have a rate smaller than 10 Hz [141]. By whittling down the incoming data to a selection that incorporates these processes, we can create a data stream that fits comfortably within our bandwidth.

A two-tiered trigger system is set in place to process collision events during data-taking and to apply filters that decide whether the event is interesting for analysis purposes before storage. Thanks to this trigger system, the data-taking rate is reduced from 40 MHz to about 1000 Hz and contains events in line with the CMS physics case. The other side of the coin is that the triggers have to be carefully designed, validated and monitored as any event that does not pass the filters is discarded forever.

In the context of the service work performed by the author, this chapter will start by providing a general description of the two tiers of the trigger system at CMS as it was designed for Run II in Section 5.3. The same section will then follow up with a more in-depth explanation of the core ideas behind the τ trigger paths. This will be followed by a description of the upgrades in the trigger system for Run III in section 5.4. With all the previous sections in hand, the way is paved to a chapter on how the τ trigger strategy was adapted for Run III in Section 5.5. Finally, a conclusion and outlook on the triggers is provided in the final section of the chapter.

5.3 Triggers in Run II

The trigger system at CMS consists of two stages that run consecutively called the L1 trigger and the HLT. Let us take a moment to put each of them into context.

5.3.1 L1 Trigger

The level 1 trigger, also called L1 trigger, is a hardware-based trigger system located at the detector. Its aim is to very coarsely process the incoming information in order to provide a first selection of interesting events and as such reduce the massive influx of collision events

to levels that are manageable by the next step in the triggering process, namely the HLT. It has a fixed latency of $3.8\ \mu\text{s}$ and reduces the incoming rate of 40 MHz to about 100 kHz. For a complete description of the hardware that constitutes the L1 trigger, the reader is referred to the Technical Design Report [142, 143]. In the scope of this work, it suffices to know that the signals from the detector are duplicated at the front-end and placed on hold in front-end buffers where it awaits the trigger decision. Meanwhile, the input signals for each subsystem are translated into so-called trigger primitives, which are objects that serve as building blocks for the next step. They are used to reconstruct the L1 trigger objects which are in turn fed into the global trigger where a final decision is made.

In the case of τ triggers, for example, calorimeter information is divided up in grids of so-called trigger towers where the energy in the ECAL and HCAL is summed. A clustering algorithm then reconstructs clusters around local energy maxima and subsequently constructs L1 τ trigger primitives from them. The isolation variables are calculated using the unused towers in the neighbourhood of the L1 τ . The L1 trigger then filters on the number of isolated objects of interest in the event. A double τ L1 trigger, for example, will then make a selection on the presence of at least two isolated L1 τ 's in the event. Because the focus of this chapter lies on the HLT, we will not go over all L1 trigger algorithms here. Instead, the interested reader is referred to Ref. [144]. It is important to note that this hardware filters only on the inputs coming from the calorimeters and muon detectors. No information from the tracker is processed at this stage.

If the event is accepted, the full event moves on to full readout so the HLT can take charge of the next step, where information related to the L1 trigger objects is often used as a seed to more complex algorithms. Otherwise, the information in the buffers is discarded. The L1 trigger menu consists of a range of L1 triggers that filter the incoming information in line with the HLT menu and thus the CMS physics case.

5.3.2 High Level Trigger

The next stage in the CMS trigger system is the HLT and it aims to further reduce the rate from 100 kHz to around 1 kHz. It is a software-based system with access to the full detector readout that runs on a cluster farm with about 24000 processor cores [141, 145, 146]. Because of the lower incoming rate, the available reconstruction time at the HLT is higher than at L1, namely around 300 ms. This allows us to perform reconstruction using techniques that are closer to the offline event reconstruction than those at the L1 trigger. In fact, the HLT software is a streamlined version of the offline reconstruction that runs in the same CMSSW software but was adapted to reduce the central processing unit (CPU) time consumption. This also means this is the first stage in data-taking where the tracker information is used and one can reconstruct charged particle tracks.

The structure of the HLT system consists of a collection of paths, each of which is a series of modules that either produce new collections of objects (such as electrons, muons, jets,

taus, etc.) or filter on particular collections. Once a filter fails, the path is terminated and the remaining modules will no longer be executed. Producers are more resource-intensive than filters and as such generally take longer to run. By alternating the usage of producers and filters, one eliminates the need to further run all following producers when a path fails. Furthermore, the modules in a path are arranged in order of complexity, meaning that modules that require more computation time are placed later on in the path. This means that the faster modules are run more often than the more time-consuming modules.

Two main differences between the offline reconstruction and the reconstruction at HLT can be found in the use of regional reconstruction and a simplification of the track reconstruction. The first refers to the limitation of detector unpacking and reconstruction to a specific region of interest. This region could be based around an L1 seed or an already existing higher-level candidate. The second point on track reconstruction is key to being able to bring the PF algorithm to the HLT, which allows for improved energy resolution of HLT objects and brings the reconstruction closer to the offline algorithm. The online tracking algorithm starts from pixel tracks that have hits in all four pixel layers and then performs two iterations where it looks for tracks in the full tracker in order of most frequent topologies – first higher p_T tracks followed by lower p_T tracks – similar to the offline algorithm. A third iteration loosens the requirement on the number of hits in the pixel tracker to three and makes use of regional tracking [147]. In 2017, an issue occurred in the pixel detector where some pixel modules became inactive. Therefore an additional recovery iteration was introduced in 2018 data-taking that uses two pixel hits as a seed if there are two overlapping non-active pixel modules in the region of the track.

The efficiency and fake rate of this tracking algorithm as a function of track p_T as measured in a simulated $t\bar{t}$ sample is shown in Figure 5.1. The efficiency is defined as the fraction of simulated particles that are matched to a reconstructed track. Figure 5.1 shows the cumulative contribution from every iteration to the final efficiency. When using all iterations, efficiencies up to 80% are observed. The fake rate is defined as the number of tracks that could not be matched to a simulated particle. Both the efficiency and fake rate were measured using a map that emulates the inactive modules. Both distributions are also shown in the case there are no issues in the detector.

The PF algorithm is used to create electron, muon, jet, missing transverse energy and hadronically decayed τ lepton objects at HLT. Often the reconstruction of these objects starts with less complex algorithms and finishes with the full reconstruction in order to reduce the timing of the paths. Electrons are first reconstructed from superclusters that are matched to pixel hits and then move to the full reconstruction. Muons are first reconstructed using only information from the muon detectors before moving to the stage where the full detector information is used. Similarly jets and transverse missing energy objects are first reconstructed using only the information from the calorimeters before moving on to the PF reconstruction. In Run II, two algorithms for b-jet tagging became available at HLT. The first is the Combined Secondary Vertex (CSV) method that was used in all paths up to 2018 data-taking. The second is the DeepCSV algorithm, which became the go-to module in 2018. Tau leptons are reconstructed

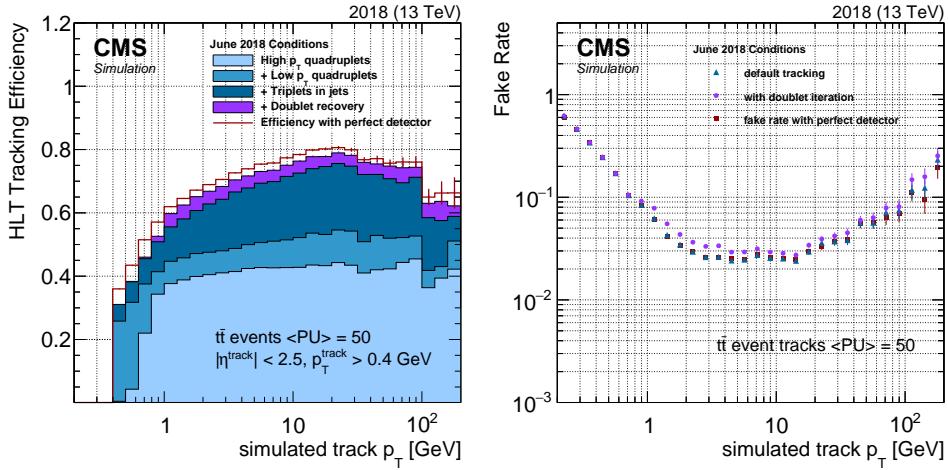


Figure 5.1: Efficiency (left) and fake rate (right) of the tracking algorithm in Run II [147].

in different stages, as will be further explained in the next section.

The HLT menu in 2018 contained around 400 paths which are grouped into different blocks corresponding to primary datasets. There are for example triggers that focus on electrons or photons that are grouped under the “EGamma” primary dataset. Any event that passes the OR of all paths within the EGamma block is saved under that primary dataset.

5.3.3 Anatomy of tau triggers

There is a multitude of paths in the HLT menu that focus on final states containing a hadronically decayed τ lepton. These triggers can be listed under different categories depending on the final state: single τ triggers, double τ triggers, $e+\tau$ triggers, $\mu+\tau$ triggers and MET + τ triggers. They all share a similar structure in their paths and go through a few stages before reaching their final filter. The general progression of all τ triggers employed in Run II is schematically represented in Figure 5.2. All paths start by filtering on the relevant L1 seeds. They then move on to two or three main reconstruction processes, each more complex but more accurate than the last, where somewhere along the line the PF algorithm is employed. Finally, a set of final filters on the most accurate set of trigger candidates from the L3 step is enforced. Let us briefly go over the different reconstruction steps below.

Level 2

The first stage of any trigger should be low in complexity in order to respect the limited timing budget and should perform a crude selection that filters out the clear cases of events that do not have the final state of interest and retains the more interesting ones for the next steps. The level 2 modules of the pure τ triggers¹ do this by only looking at energy deposits in the calorimeter with no input from the trackers. In order to save time, the reconstruction of L2 τ_h ’s is done regionally around the L1 τ_h ’s using the anti- k_T algorithm with a distance parameter of 0.2. A

¹Pure τ triggers refer to triggers that only filter on the presence of τ_h objects such as a single τ or a double τ trigger.

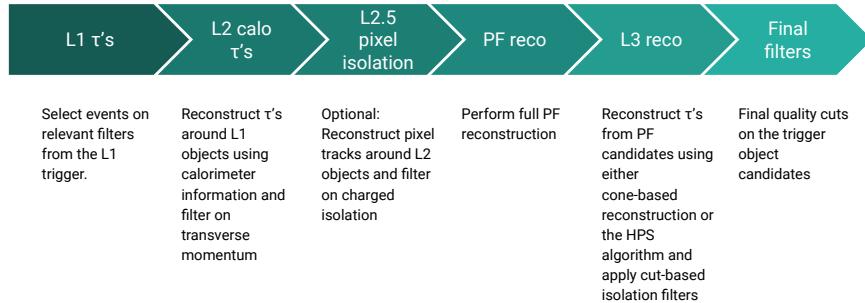


Figure 5.2: The general progression of a τ HLT path. The L2.5 step is only applicable to double τ paths.

filter module that selects on the size of the transverse momentum of the L2 candidates makes the final decision to move on to the next step or remove the event. For cross-triggers such as the $e+\tau$ and $\mu+\tau$ triggers the L2 step will instead rely on the second, non- τ object for the first set of filters.

Level 2.5

Full reconstruction as described in the level 3 stage is quite time-consuming for τ_h objects, being high level objects that need to be built up from all their decay products. Whereas $e+\tau$, $\mu+\tau$ and $\text{MET}+\tau$ triggers can rely on the second object in its final state to perform initial filters, the double τ triggers would immediately start the full reconstruction of not one but at least two hadronically decayed τ 's. In order to decrease the average timing of the path, a third stage is introduced in the double τ and single τ triggers to filter out failing events before the full PF reconstruction takes place.

The algorithm uses both the L2 τ_h 's and pixel detector information as input. Tracks are formed from hits in a region around the L2 τ_h 's which are subsequently used to form vertices. If no vertex is found, the path moves on to the L3 stage. In case at least one vertex is found, the vertex with the largest sum of squared transverse momenta is chosen as the primary vertex. A charged isolation is computed by summing the transverse momentum of all tracks that are found in the neighbourhood of the τ_h candidate and that have at least 3 hits. Finally, a filter on this charged isolation terminates the path for that event if this quantity is higher than a certain threshold.

Level 3

In the final steps of the τ triggers, hadronically decayed τ leptons are reconstructed using the PF candidates from their decay products. For most of the τ paths, τ_h 's are reconstructed globally. The only exceptions to this are the double τ paths, which are CPU-consuming triggers, where the reconstruction is performed regionally around L2 τ_h 's. Up until 2018, all τ -related paths used a cone-based reconstruction [148] for building the τ_h objects. After that, the double τ , $e+\tau$ and $\mu+\tau$ triggers moved to the HPS algorithm, which is more in line with offline reconstruction and allows for the separation of the τ_h objects according to their decay mode. Once the τ_h candidates are reconstructed, their isolation is computed as the sum of all charged hadrons and photons that are not part of the τ_h candidate and are found within $\Delta R = 0.5$ of the leading charged hadron. Finally a filter is applied on the number of such isolated τ_h 's.

5.4 Triggers in Run III

With the Run III data-taking period, the LHC aims to go beyond Run II by doubling the integrated luminosity. This is done by aiming for a high peak luminosity and maintaining it for hours on end. At the same time the centre-of-mass energy is bumped up to 13.6 TeV and the mean pile-up is expected to be higher compared to Run II [149]. In order to be up for the challenge of optimally processing this data at a target output rate of 1.3 kHz, CMS had to update its trigger strategy [150].

One of the main upgrades that also has the High-Luminosity upgrade after Run III in mind comes with the addition of graphical processing units (GPUs) to the HLT farm. Every node has received an NVIDIA GPU to offload some of its tasks to and speed up processing. In order to prepare the software for this addition, the Patatrack project [151] was launched. Its main purpose was to create algorithms that could perform pixel tracking and vertexing on GPUs. Simultaneously some of the algorithms pertaining to calorimeter reconstruction were updated for GPU compatibility as well. If an event takes place and gets to an HLT farm node, the software will check if the GPU is available and run the track and vertex reconstruction on there. This technique allows for great time reduction and to employ these pixel tracks at earlier stages. These new features come with some improvements in the tracking algorithms as well. Instead of the multiple iterations as was used in Run II, this time a single global iteration is employed that starts from Patatrack pixel tracks that have at least three pixel hits. The efficiency of this new algorithm compared to the one we had in Run II can be seen in Figure 5.3.

Further improvements can be found in the L1 sector, an overview of which is given in Ref. [149]. It is interesting to note that due to upgrades in the HCAL electronics and the addition of new variables in the muon sector, Run III will herald the first time that the CMS L1 algorithms can identify long-lived particles. HLT paths targetting long-lived signatures were developed accordingly. These paths benefit from improved muon and τ_h reconstruction

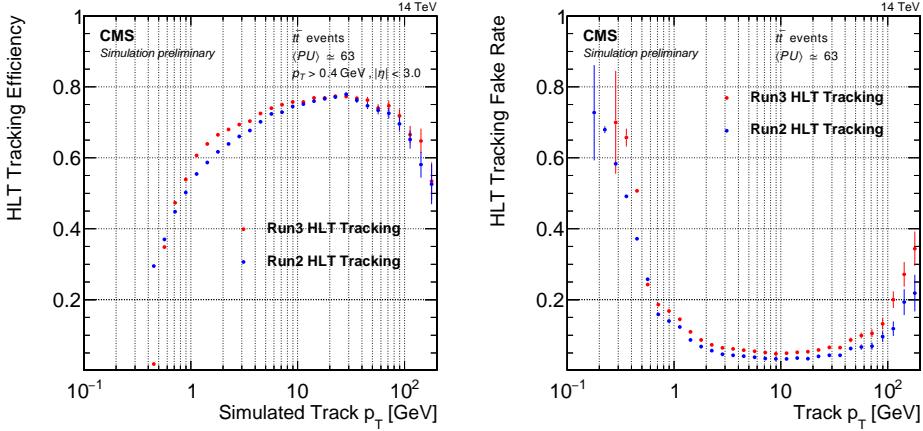


Figure 5.3: Efficiency (left) and fake rate (right) of the tracking algorithm in Run II and Run III [150, 152].

as well as b-tagging with the use of Patatracks.

5.5 Tau Triggers in Run III

In preparation for Run III, we took a look at the existing τ paths and established the areas where we could make some improvements. It was found that the greatest advancements could be made in the reconstruction techniques and filters of the L2, L2.5 and L3 stages by making use of the new GPU infrastructure as well as porting offline identification techniques to the HLT. These come to life in the two main projects of the τ HLT for Run III developments, namely an improved L2 sequence and the inclusion of the DeepTau@HLT algorithm. The next two subsections will be dedicated to them. A third subsection will then go over the smaller upgrades and what the final menu in terms of τ triggers looks like along with its performance.

5.5.1 Improved L2 sequence

As mentioned earlier, the level 2 and 2.5 sequences were introduced in pure τ paths to reduce the impact of the large time cost of the full τ_h reconstruction. They are not designed, however, to withstand the increased instantaneous luminosity and pileup of Run III. Furthermore, they do not make use of the new patatracks in an optimal way. Additionally, since the TSG group has found the CaloTowers to be obsolete for Run III, the dependency of the L2 sequence on them is no longer desired. To address these issues, maintain a high efficiency and keep the timing under control, a new machine learning strategy to tag hadronic τ_h in the early stages of the HLT was designed.

From a technical point of view, this new machine learning technique is based on a convolutional neural network (CNN) [153] that was designed, trained and validated by Valeria D'Amante, who was the driving person behind the improved L2 sequence. We will refer to this CNN as the L2NN. CNNs work exceptionally well on image recognition, which is something we can adapt to τ 's in the CMS detector by translating their components into an image grid around the seed. The idea is to remove the modules and sequences that perform the L2 and L2.5

reconstruction and filtering and replace them with a few modules that calculate L2NN scores on which we can filter accordingly. The updated algorithm starts from the L1 τ_h candidates and gathers patatrack and calorimeter information within a cone of $\Delta R < 0.5$ around them. This information is divided into a 5×5 grid in the $\eta \times \phi$ plane around the respective L1 candidate where each cell has a size of 0.2×0.2 . These information grids, along with the number of vertices and the L1 candidate observables, then serve as input variables to the L2NN. As such there is no need for τ_h reconstruction in the early stages and one only uses existing input variables.

An L2NN output score is calculated for every L1 candidate, which is accordingly marked as “good” if the score is higher than a path-dependent L2NN score threshold or has a large p_T (higher than 250 GeV). The event then passes the L2 stage if there are enough “good” candidates according to the specific path needs, i.e. two good candidates in the double τ path. The exact value of the L2NN threshold is determined by measuring at which value the rate of this sequence is equal to a set target rate. As this depends on the data-taking conditions, the threshold needs might vary over time. At the start of Run 3 however, the L2 target rates were set to 5.4 kHz for the double τ path and 1.9 kHz for the single τ path.

The efficiency² of the new method compared to the cut-based sequence can be found in Figure 5.4. It can be seen that for the double τ trigger, the efficiency of the new L2NN sequence is equal to or higher than the one obtained using the cut-based strategy. There is some efficiency loss at high pseudorapidity, which is probably due to the low statistics in this phase space for training. The new single τ sequence sees a slight drop in efficiency compared to the cut-based sequence that can be traced back to the fact that the neural net was only trained on double τ events. Future iterations will make sure to perform a separate training. The main advantage of the new L2 sequence comes in a great timing reduction of 9 ms (or 39%) in the single τ trigger.

5.5.2 DeepTau@HLT

For offline τ_h identification, a neural network called DeepTau was developed by Konstantin Androsov et al. during Run II [154] to improve on the existing identification methods. It is a deep neural network (DNN)-based algorithm that makes use of high-level and low-level inputs to simultaneously discriminate genuine hadronically decayed τ 's against jets, electrons and muons. One of its main features is the convolutional layers processing the activity in the detector in the neighbourhood of τ_h signatures. Starting from τ_h candidates reconstructed by the HPS algorithm, one collects information from reconstructed particles close to the τ_h axis and translates them into two overlapping grids in $\eta \times \phi$: a fine-grained inner grid of 11×11 cells with size 0.02×0.02 and a coarser outer grid of 21×21 cells with size 0.05×0.05 .

²The efficiency is defined as the fraction of events with at least one (two) τ_h ('s) that pass the L2 threshold of the single τ (double τ) trigger.

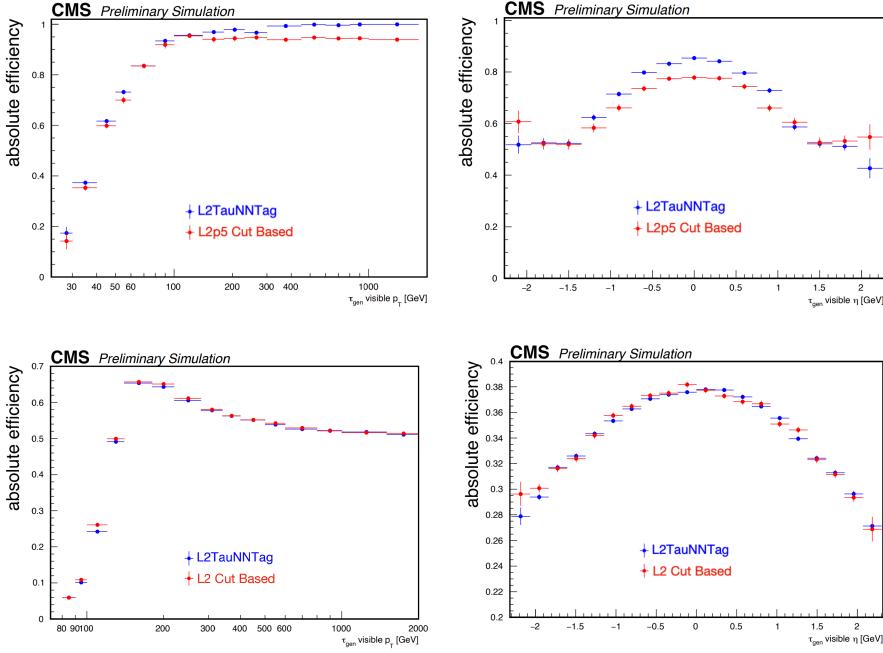


Figure 5.4: Efficiency of events with one (two) τ_h ('s) to pass the new L2 requirement of the relevant sequence as a function of visible p_T (left) and visible η (right). Shown on the top for the double τ sequence and on the bottom for the single τ sequence. Plots made by Valeria D'Amante.

Each cell contains for each of the types³ of object the relevant information of the candidate of that type with the highest p_T present in the cell. Together with high-level variables such as the τ_h four-momentum and charge, isolation variables, properties of the secondary vertex if applicable, and observables related to the energy distributions in the strips, these serve as inputs to the neural network.

An overview of the architecture of the DeepTau network is shown in Figure 5.5. There are three subnetworks, each handling a different type of input: One for the high-level variables, one for the inner grid and one for the outer grid. A final set of layers uses the outputs produced as a result of these subnetworks and produces four outputs corresponding to genuine taus, jets, muons and electrons. The subnetwork for the high-level variables consists of three fully connected layers, the outputs of which are fed directly into the final network. The inner and outer grid subnetworks first process the input variables through a number of fully connected layers in order to reduce the dimensionality of each cell. The 64 output variables in each cell are then used in convolutional layers, the output of which goes to the final layers.

³There are seven types of objects in CMS that are relevant to the DeepTau network. These consist of PF charged hadrons, PF neutral hadrons, electrons, PF electrons, PF photons, muons and PF muons.

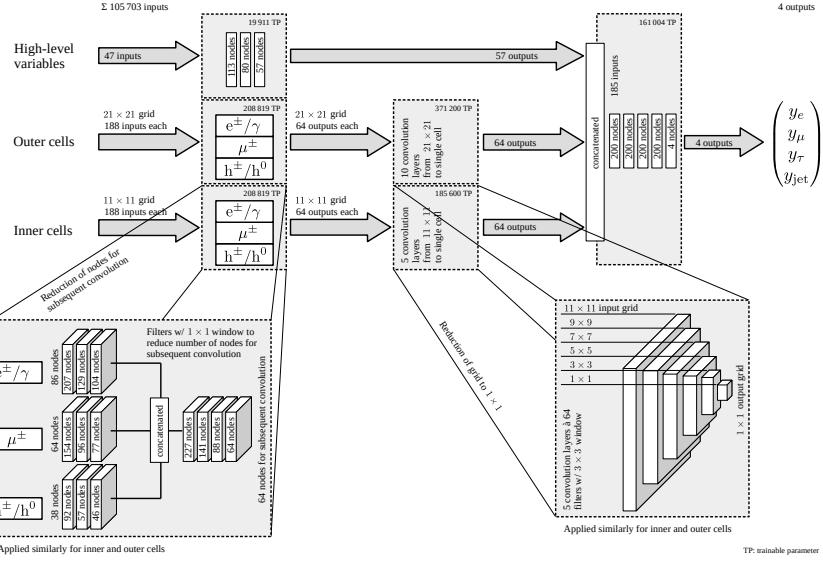


Figure 5.5: Architecture of the DeepTau neural network [154].

This discriminator was integrated into the τ -related HLT paths in preparation for Run III with a focus on the discrimination against jets. As the HLT menu runs in the same CMSSW framework as offline reconstruction, the code for evaluating the neural network was already available. However, since offline and online processing uses different, incompatible input collections, the code in question had to be adapted to work with all input collections and to compute any input variable to the DNN that would not be readily available in the online collections. The main change is in the electron and muon input components for the inner and outer grid. The collections that go into the neural network are exchanged for empty collections due to a lack of out-of-the-box availability of all the necessary components, namely fully reconstructed electrons and muons. If the dependency on these variables would be kept, it would mean triggers like the double τ trigger would also need to run additional modules to make sure electrons and muons are available. This is clearly an undesired dependency to have in an environment where processing time is limited. At the time of writing, the training used at the HLT is the same as the offline training although there are plans in the making to perform a retraining dedicated to the trigger system. Nonetheless, the performance of this discriminator is greatly improved over the previous cut-based sequence as shown in Figure 5.6. These show the efficiency and rate of HLT paths that make use of the DeepTau discriminator at L3.

Once the technical implementation in CMSSW and the first validations were performed, we could move on to defining DeepTau working points for actual use in the HLT paths. As the rate at high $\tau_h p_T$ is significantly lower than the rate at low p_T , it was chosen to make these working points p_T -dependent. This allows to maximise efficiency in the high p_T spectrum where there are fewer incoming events while at the same time reducing the rate at low p_T . The exact function of the threshold is obtained by maximising the algorithmic efficiency of the L3 sequence while targeting a specific rate. For all trigger paths, we started from a collection of linearly decreasing functions as illustrated in Figure 5.7 where the parameters t_0 , t_1 , and t_2

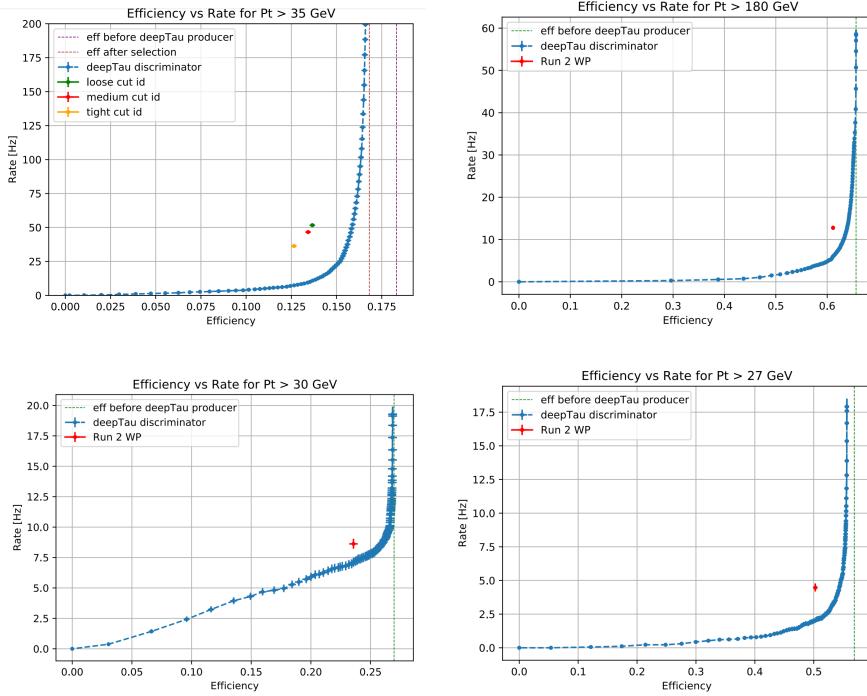


Figure 5.6: Efficiency versus rate for a double τ trigger (top left), a single τ trigger (top right), an $e+\tau$ trigger (bottom left) and a $\mu+\tau$ trigger (bottom right). Plots made by Anna Mascellani.

are optimised. The target rates used for this process were kept the same as in the 2018 trigger menu but were scaled up to an instantaneous luminosity of $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

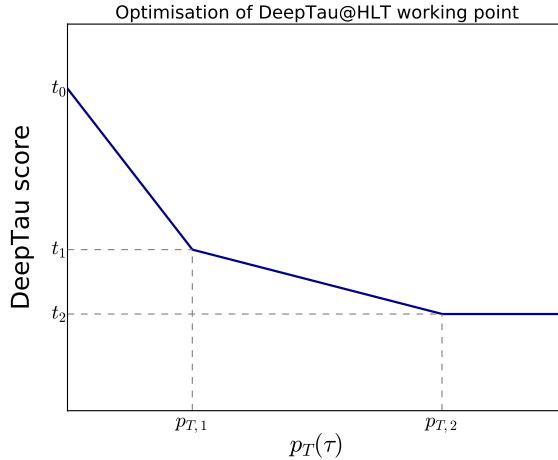


Figure 5.7: Example of the type of DeepTau score vs p_T function used to create the p_T -dependent working points for the DeepTau@HLT sequence. For set values of p_T , the values for t_0 , t_1 and t_2 are optimised in terms of efficiency for a certain target rate. Figure by Anna Mascellani.

5.5.3 Final strategy and performance

With these two improvements in place, we can put together the final paths for the Run III menu. The general structure of τ paths in Run III is shown in Figure 5.8, where the positions

of all updated sequences are highlighted. As you can see the L2 and L2.5 sequences are taken out and replaced by one improved L2 sequence and the L3 sequence remains but now uses the DeepTau module for isolation requirements. In Run II paths the L3 sequence was followed by a quality cut on L1 matching requirements of the selected τ_h 's. In order to reduce the timing of the paths, this module was adapted to fit in front of the DeepTau sequence so only L1-matched τ 's get their DeepTau score calculated, a process which can be rather time-consuming.



Figure 5.8: The general progression of a Tau HLT path. The L2.5 step is only applicable to ditau paths.

Two other areas where improvements were made are found in the reconstruction areas. In the 2018 menu, the L3 sequence used cone-based reconstruction in some paths and the HPS algorithm in others. All τ trigger paths intended to be active in Run III have now been unified to make use of the HPS algorithm. This is not only a desired update in general but also a required one to be able to run the DeepTau algorithm. Unlike the cone-based reconstruction, the HPS reconstruction provides the candidates with decay mode information which serves as input to the DNN. Secondly, all paths have been updated to use global reconstruction instead of regional reconstruction. In the past, it was shown that for the double τ path in particular, this brought about large increases in timing (of the order of 10 ms) but it was found that with the GPU-based tracking in place, this timing increase is not as severe. A disclaimer should be placed here that all aspects discussed in this subsection, but especially the global reconstruction, can be subject to changes during the Run III data-taking to meet timing and rate needs of the menu depending on what is observed in the DQM. Everything described here is how the first menu was designed to be used at the start of the run.

The final four proposed τ trigger path types then consisted of a double τ path, a single τ path, a $\mu+\tau$ path and an $e+\tau$ path. The observation was made that in 2018 the pure rate⁴ of the τ +MET trigger was negligible and the phase space where the trigger reached useful efficiencies

⁴The pure rate refers to the rate of data-taking that is not covered by any other trigger and purely stems from passing the trigger in question.

was almost fully covered by the single τ triggers. In order to reduce the timing contribution of τ triggers to the total timing, it was decided to therefore remove all $\tau + \text{MET}$ triggers from the menu.

For validation purposes, every τ trigger came with a monitoring trigger. They consist of a $\mu + \tau$ trigger where the τ parts are identical to the parent trigger. This way, a tag-and-probe method can be used later on where the muon serves as a tag and the τ leg serves as the probe. Lastly, it should be noted that as a contingency the old paths using the old L2 and L2.5 stages as well as cut-based isolation in the L3 stage were updated to Run III standards and retained in the menu. They are disabled by default, however, meaning that in normal conditions these paths will never be run. In case it is observed in DQM that despite the thorough validation in development something goes wrong with the machine learning sequences during data-taking, one can exchange them with the old paths that are tried and tested by Run II for the remainder of the run.

A comparison of the total efficiency of the τ leg in the $e + \tau$, $\mu + \tau$ and double τ paths as used in Run III and the Run II version of the path can be found in Figure 5.9. These were measured using the monitoring triggers in a tag-and-probe method. For all p_T values, an increase in efficiency of the total path is observed and in the plateau, the efficiency is around 86% for the double τ path. Similar observations are made in the cross triggers where the efficiency in the plateau reaches 90%. Note that while the $e + \tau$ and $\mu + \tau$ triggers use the same sequences for the τ section of the paths, they do not necessarily have the same efficiency as there are differences in L1 seeds, p_T thresholds and L2NN and DeepTau score thresholds.

5.6 Conclusions

In this chapter, we discussed the CMS trigger system with a focus on the HLT. While this overview was rather general, a more in-depth discussion on the structure of τ paths took place. This was followed by the evolution of the trigger system and τ paths into Run III. We saw that the inclusion of machine learning techniques allowed for high-efficiency triggers that operate at similar or better time scales and rates. This will continue to evolve with preparations for a dedicated DeepTau training at HLT in progress as we speak. Also, the L2 CNN keeps evolving and carefully crafted plans are being made for a new training that will deal with some of the open issues such as efficiency loss at high pseudorapidity and separate trainings for different trigger types.

This will continue to evolve in data-taking eras after Run III. Dedicated teams are already designing paths for Phase II, which will follow after Run III, where the detector is upgraded with new subdetectors. They are preparing the paths to incorporate these changes and are updating the DeepTau module for that purpose. During Phase-II, tracking information will already be available at L1, which will open new doors for the trigger menu [155]. With great improvements in areas other than τ 's already implemented or lined up, we keep pushing the boundaries of the system into better menus that allow for tougher but greater data-taking

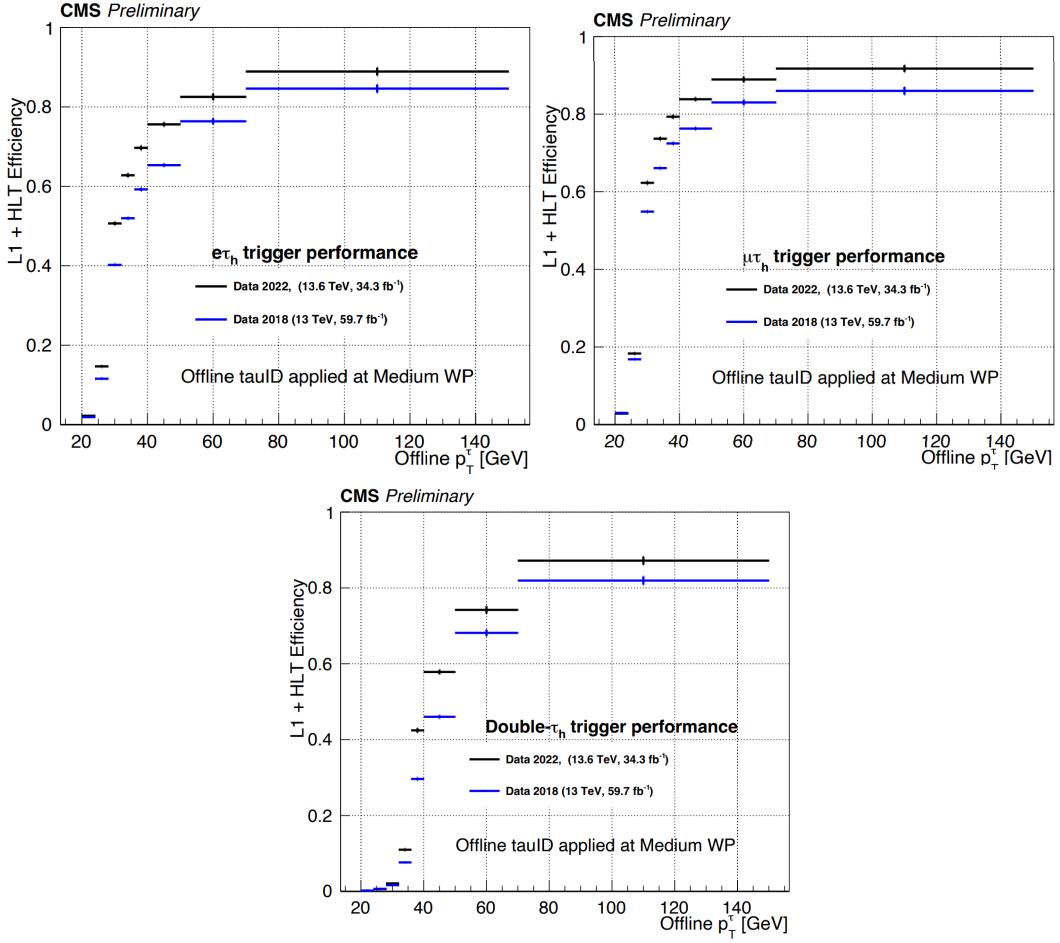


Figure 5.9: Comparison of the total efficiency of the $e\tau_h$ path (top left), $\mu\tau_h$ path (top right) and $\tau_h\tau_h$ (bottom) path as a function of offline $\tau_h p_T$ in Run II (blue) and the Run III (black) working point [8].

conditions as well as difficult candidates such as displaced objects.

Personal contributions

My contributions to τ triggers had already started before I became τ trigger convener. This came in the form of service work with purely technical tasks. I started by optimising the fitting procedure used in calculating scale factors for all τ trigger paths. This was followed by the actual measurement and validation of full Run II τ trigger scale factors. Not long after, I performed the study of factorisation in the double τ triggers, found there to be some issues with it, designed a procedure to obtain corrections to address these issues and measured these corrections.

When I became a τ trigger convener my tasks were twofold. On one hand, I had managerial roles. I was in charge of the development of τ triggers for Run III and supervised the students working on the different subprojects. I followed all trigger meetings and gave regular reports there as well as chaired the τ trigger meetings where all important information from the general meetings was propagated by me. At the same time, I supervised

the DQM tasks and was responsible for the bookkeeping of service tasks in the group.

On the other hand, I kept performing technical work: development of the code to calculate Deep-Tau scores at HLT and the implementation in the CMSSW framework, design of the new τ trigger paths for Run III that included all new updates and their implementation in the trigger menu, implementation of new updates in triggers that included τ 's but were not part of our group, updates of trigger bits in the offline framework, and timing and rate measurements of the new paths.

6

A search for supersymmetry in a multilepton final state

6.1 Personal foreword

After wrapping up my traineeship at CERN and returning to Ghent in the summer of 2017, I sat together with Professor D. Dobur to discuss a subject for my master’s thesis. Searches for SUSY came up and as my younger self had a bit of a fascination with the theory, I eagerly accepted. I had come into contact with the theory side of this extension to the SM in lectures but I had very little feeling with the experimental side. For the next academic year, I worked together with Willem Verbeke, who was experienced in searches for signs of the presence of this additional symmetry in data taken with the CMS detector. Here I was focusing on probing new variables that have good discriminating power between signal and background and that could be used by Willem in his analysis. While the original idea did not work out, it was a great learning experience not only regarding SUSY in practice but regarding searches at experiments at the LHC in general. The remnants of this work were archived in my master’s thesis [55].

When the time arrived to write a proposal for PhD funding, I once again sat together with Professor D. Dobur. I preferred to perform a search for new physics again, but this time in a different area than SUSY. While we decided to write a proposal on a search for HNLs and their coupling to tau neutrinos, which will be the subject of the next chapter, we also agreed that I would first work alongside Willem Verbeke for a year while he was working on a new search for SUSY, this time using Run II data, to learn more on the structure of an analysis and to get insight into different analysis techniques. Within the context of the work, I focused on the hadronic tau sections, as they were most relevant for the following analysis. More specifically, I worked mostly on finding the optimal τ_h identification and subsequently on background estimation of nonprompt τ_h .

Because I was not the main analyser of this search and focused on only some aspects of the analysis, I will refrain from giving a full description of the search and all its search channels.

The intention of this chapter is not to illustrate the SUSY search from A to Z where we go through all signal models with their motivation and the complete analysis strategy but to highlight the work I have performed in the first year of my contract. Therefore, I will instead stick to a subset of signal processes and channels that support that work. However, as there is a large overlap in the methodology used in this analysis and the one presented in the next chapter, I will take the time in some places to give more detailed explanations of the concepts and techniques behind these overlapping aspects. This will be the case in the object selection, nonprompt light lepton background and systematics sections.

6.2 Introduction

Chapter 3 introduced SUSY as an extension to the SM that inserts an additional symmetry between fermions and bosons. The existence of such a supersymmetry comes with the implication that there are superpartners for all the SM particles, which in turn means we can target these new particles in searches. The CMS experiment has been active in such SUSY searches since the first years of data-taking and hunts for SUSY signatures in a wide range of final states. So far no sign of new superpartner particles has been observed.

One of these analyses in the electroweak production category is described in Ref. [7]. It explores the direct production of charginos and neutralinos in processes with multilepton final states using a dataset of proton-proton collision events recorded by the CMS detector during the Run II data-taking period. To accomplish this, it applies a strategy that uses a mix of cut-based search regions and parametric neural networks. This chapter describes a subset of that analysis that focuses on two varieties of a slepton-mediated decay model that favour decays to τ leptons. In this context, events with three leptons of which at least one is a τ_h are investigated. While it is true that events with three light leptons will also contribute slightly to the final sensitivity to these models, we will not consider them further in this chapter. R-parity conservation is assumed in the analysis, meaning at least two LSPs will be present in the final state. They will not interact further with the detector and lead to missing transverse energy signatures.

The structure of this chapter is as follows. The foundation of the analysis, namely the signal models under consideration, will be highlighted in Section 6.3. This will be followed by a few words on the data and simulation in Section 6.4. We then move on to an explanation on object identification and selection in Section 6.5 after which the analysis strategy and event selection are presented in Section 6.6. All elements that go into the background prediction are laid out in Section 6.7. The last remaining ingredient at that point is the systematic uncertainties that creep into the analysis, an overview of which is provided in Section 6.8. Finally, we are ready to look at the results and interpretation in Sections 6.9 and 6.10.

The details of this analysis were transcribed in the analysis note linked to it. This chapter is heavily inspired by that note and the paper, both in content and in structure [7].

6.3 Signal Models

Due to the high number of free parameters in the MSSM, we make use of simplified SUSY models in this analysis. The only free parameters in these simplified models are the masses of a handful of sparticles [52–54]. The full analysis considers three such models of which one, referred to as slepton-mediated decay, is described below.

All slepton-mediated models in this analysis assume wino-like $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^\pm$ – meaning they are mass-degenerate mixtures of the charged and neutral winos – that have the same mass while the LSP is assumed to be bino-like, meaning that $\tilde{\chi}_1^0 \approx \tilde{B}$.¹ In scenarios where sleptons have a mass that is settled between the mass of the chargino $\tilde{\chi}_1^\pm$ and the mass of the LSP $\tilde{\chi}_1^0$, sleptons can act as the mediator of the decay of the original chargino-neutralino pair. The diagrams considered for this type of process are shown in Figure 6.1. The chargino can decay to a lepton and a sneutrino, which will further decay to a neutrino and an LSP. Alternatively, the chargino can also decay to a neutrino and a slepton. In that case, the slepton will in turn decay into a SM lepton and the LSP. The heavy neutralino $\tilde{\chi}_2^0$ would decay to either a slepton and a lepton or a sneutrino and a neutrino. The final state of the last option would lead to events with a single lepton so we will not consider it further.

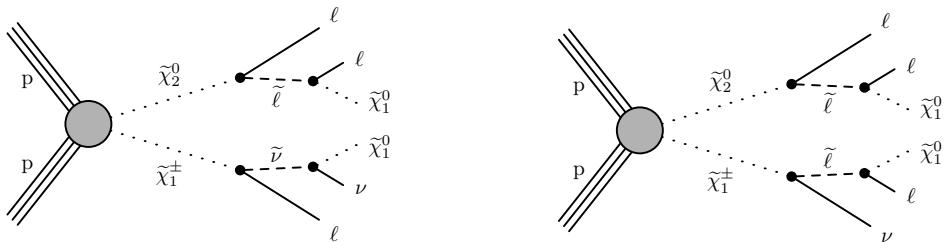


Figure 6.1: Diagrams depicting processes where a chargino-neutralino pair is produced and subsequently decays via a slepton and a sneutrino (left) or via sleptons (right) [7].

Within the slepton mediated models regarded in this analysis, there are multiple different submodels hypothesizing the type of slepton that mediates the decay. Scenarios where the chargino only couples to right-handed sleptons while the neutralino decay can still occur through left-handed sleptons, is called the τ -enriched scenario. The origin behind the name is that in this scenario, right-handed sleptons would only couple to the Higgsino component of the chargino. This results in a preferred decay to τ leptons since they have the highest of the three lepton masses and thus the largest Yukawa coupling. If light flavour sleptons have such a high mass that they are decoupled, it would only be possible for the chargino and neutralino to decay through τ sleptons, implying an exclusive decay to τ leptons. This last scenario is called the τ -dominated scenario.

¹Winos and binos are the superpartners of the electroweak bosons in the SM before spontaneous symmetry breaking.

The kinematics of the leptons in the event as well as the amount of missing transverse energy will depend on two mass differences: the one between the original chargino/neutralino and the LSP and the one between the slepton mediating the decay and the LSP. The first-mentioned of these mass differences will be further denoted as $\delta m = m_{\tilde{\chi}_2^0} - m_{\tilde{\chi}_1^0}$ and will have an impact on the energy available to the leptons in the final state. The mass difference between the sleptons and the LSP will have an impact on the relative distribution of the available energy between the leptons in the final state and will be further characterized by a parameter x , the definition of which can be found in equation 6.1.

$$m_{\tilde{l}} = x \cdot m_{\tilde{\chi}_2^0} + (1 - x) \cdot m_{\tilde{\chi}_1^0} \quad (6.1)$$

Three values for x are studied in the analysis:

- $x = 0.05$: The mass of the slepton is close to the LSP mass. This means that there is little energy left for the lepton produced in the decay of the sleptons which means that they typically will be soft.
- $x = 0.5$: With a mass exactly halfway between $m_{\tilde{\chi}_2^0}$ and $m_{\tilde{\chi}_1^0}$, the phase space available to each lepton will be the same due to the identical mass difference at each stage of the process. This means the momentum spectrum of all leptons will be similar.
- $x = 0.95$: This situation is comparable to the first scenario but this time the mass of the slepton is close to $m_{\tilde{\chi}_2^0}$. The leptons produced in the decays of $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^\pm$ will have a low transverse momentum spectrum due to the compressed phase space.

6.4 Data and event simulation

6.4.1 Collision data

Three sets of proton-proton collision data at $\sqrt{s} = 13$ TeV, corresponding to three years of Run II data-taking, are scrutinized in this search. They consist of 35.6 fb^{-1} of 2016 data, 41.5 fb^{-1} of 2017 data and 59.7 fb^{-1} of 2018 data for a total integrated luminosity of 137.2 fb^{-1} , collectively referred to as Run II. As running conditions were slightly different across the three years, they are treated separately in the analysis with each having their own set of simulations, calibrations and validations and forming separate inputs to the final fit. The analysis uses events from the single-electron, single-muon, double-electron, double-muon and electron-muon primary datasets. These have been cross-cleaned to remove any potential double counting.

6.4.2 Event simulation

All simulation samples follow the same structure of the datasets, with one set of simulations per year of data-taking. The signal samples for all models described in the last section have been generated with the `MADGRAPH5_AMC@NLO` program [156, 157] at LO in the strong coupling constant and with up to two additional partons in the matrix element computations. Background process samples for WZ production and ZZ production through the $q\bar{q} \rightarrow ZZ$

and $H \rightarrow ZZ$ channels were generated at next-to-leading order (NLO) using the POWHEG v2 generator [158–161] while the ZZ process in the production mode $gg \rightarrow ZZ$ was produced at NLO using the MCFM v7.0 generator [162]. With the exception of background processes that consist of two top quarks along with two massive bosons, which were generated at LO, all other processes were generated at NLO with the MADGRAPH5_aMC@NLO program.

6.5 Object selection and identification

All particle candidate objects used in the analysis are reconstructed using the PF algorithm discussed in Section 4.3.3. Not all of the reconstructed lepton candidates are useful to the analysis, however. They could be fake, for example from misreconstructed hadrons or photon conversions in jets, or they could be genuine leptons that originate from sources that are not of interest to the signal process, such as leptonic b-jet decay. In general, we refer to leptons that originate from sources of interest – i.e. the decay of electroweak bosons or superpartners – as *prompt* leptons and leptons from all other sources as *nonprompt*. In order to make sure that we select prompt leptons as efficiently as possible, additional identification criteria are applied for the different lepton flavour candidates.

In general, there are three working points of identification criteria for the object candidates called *loose*, *medium* and *tight*, each being more strict in their selection than the last. The *loose working point* is designed as a baseline selection that rejects most misidentified jets and leptons while maintaining a high efficiency for prompt objects. For the object to be regarded in any step in the analysis, it has to pass the loose working point. The *medium* or *fakeable object (FO)* working point is designed with nonprompt background estimation in mind and is built to select a good portion of nonprompt objects along with the prompt ones. The *tight working point* is regarded as the final working point and is intended to select prompt leptons as purely as possible. All objects selected in the final selection need to pass this working point.

6.5.1 Muons

The **loose baseline selection** is built on a few core principles regarding impact parameters and isolation. Firstly, loose muons must be observed within the muon system, which spans up to $|\eta| < 2.4$, and must have a p_T higher than 5 GeV. Secondly, the track must have small impact parameters² in the xy -plane $|d_{xy}|$ and in the z -direction $|d_z|$. Also, the *significant impact parameter (SIP)* value in all three directions, defined as the impact parameter in 3D-space divided by the uncertainty on this value, must be lower than the threshold of 8. This selection helps to reduce the contributions of genuine nonprompt leptons from i.e. b-quarks, which usually have a high impact parameter due to the lifetime of their parent.

Thirdly, we expect little hadronic activity in the neighbourhood of a genuine, prompt muon. Misidentified muons from jets or muons from heavy quark decay will typically be surrounded

²The impact parameter is defined as the distance between the primary vertex and the point of the track closest to the primary vertex.

by the other hadrons. Therefore gathering information on the crowdedness in the vicinity of muon candidates by taking the scalar sum of the p_T of all PF candidates in a cone around the muon can offer valuable information on its origin. This value is called the (*absolute*) *isolation* I_{abs} and is defined as

$$I_{\text{abs}} = \sum_{\text{charged}} p_T^{\text{charged}} + \max \left[0, \sum_{\text{neutral hadrons}} p_T^{\text{neutral hadrons}} + \sum_{\gamma} p_T^{\gamma} - p_T^{\text{neutral pileup}} \right]. \quad (6.2)$$

Here the contribution due to charged particles requires them to originate from the primary vertex in order to reject pileup contributions from other vertices. For the neutral particles, where this track information is not available, the contribution from neutral pileup ($p_T^{\text{neutral pileup}}$) is predicted using the effective area technique described in Ref. [163] and is subtracted from the neutral contribution to the isolation. In this analysis, the isolation variable of choice is called *mini-isolation*, which is defined by the fact that the size of the cone will vary with muon p_T following the definition:

$$\Delta R(p_T) = \frac{10 \text{ GeV}}{\min(\max(p_T, 50 \text{ GeV}), 200 \text{ GeV})}. \quad (6.3)$$

The advantage of using a variable cone size is that it takes into account that the higher p_T jets in which the muons are embedded are more collimated. By using a smaller cone size for higher- p_T muons, we reduce contamination of the isolation sum by sources other than the jet. The final variable for isolation used to select loose muons is the *relative mini-isolation*, which is defined as the mini-isolation divided by the reconstructed muon p_T .

Identification requirements are placed on the matching between the muon segments in the tracker and the muon system. These are referred to as the *PF-medium* requirements, more details on which can be found in reference [117].

In order for a muon to pass the **FO working point**, it should pass the loose working point along with an increase in the aforementioned p_T cut to 10 GeV. Furthermore, the muon can either pass the same lepton multivariate analysis (MVA) cut as for the tight working point or, if that fails, it should at least pass two additional cuts. The first of these is a requirement on the p_T^{ratio} , which is defined as the ratio between the muon p_T and the p_T of the closest jet. In case a muon is part of a jet, this ratio will be relatively low as the jet p_T is distributed across all its constituents. Secondly, there is a selection on the DEEPFLAVOR b-tagging score of the closest jet to discriminate against leptons that originate from heavy quark decays. In the case of muons, the threshold on this latter cut is made to be a sliding cut that adapts to the p_T of the lepton.

Tight muons are required to pass the loose and FO analysis working points as well. To increase the discrimination power between prompt and nonprompt muons, a cut on the output score of a gradient boosted decision tree (BDT) is applied. This BDT was trained to operate for both muons and electrons [7, 164, 165]. Among the input variables are the properties of the jet that contains the light lepton such as its DEEPFLAVOR b-tagging score [166], the ratio of the lepton p_T to the jet p_T and the component of the jet momentum transverse to the direction of the lepton, as well as impact parameter variables and isolation variables. With

this tight selection, an identification efficiency for prompt muons running between 90% and 99% is obtained with a 5 to 10% misidentification rate. Here, the identification efficiency (misidentification rate) is defined as the fraction of prompt (nonprompt) muons passing the baseline selection that also passes the BDT working point. A summary containing the exact definitions of all muon working points can be found in Table 6.1.

Table 6.1: Selection requirements that define the muon ID WPs. Cuts that are labelled with † are applied in the case that the muon does not meet the lepton MVA requirement. The sliding cut on $\text{DEEPFLAVOR}^{\text{closest}}$ is defined as a linear interpolation between two thresholds as a function of the muon p_T . This linear interpolation runs from a threshold on $\text{DEEPFLAVOR}^{\text{closest}}$ of $\frac{\text{Medium WP}}{2}$ at 10 GeV to $\frac{\text{Loose WP}}{5}$ at 50 GeV in 2017 and 2018 data. In 2016 data these thresholds take on the values $\frac{\text{Loose WP}}{1.5}$ and $\frac{\text{Loose WP}}{5}$ at 10 and 50 GeV respectively.

| Cut | Loose | Fakeable Object (FO) | Tight |
|--------------------------------------|-------|--------------------------|---------|
| $ \eta < 2.4$ | ✓ | ✓ | ✓ |
| p_T (GeV) | > 5 | > 10 | > 10 |
| $ d_{xy} < 0.05$ cm | ✓ | ✓ | ✓ |
| $ d_z < 0.1$ cm | ✓ | ✓ | ✓ |
| $\text{SIP}_{3D} < 8$ | ✓ | ✓ | ✓ |
| $I_{\text{mini}}^{\text{rel}} < 0.4$ | ✓ | ✓ | ✓ |
| PF-medium | ✓ | ✓ | ✓ |
| lepton MVA | — | — | > 0.4 |
| $\text{DEEPFLAVOR}^{\text{closest}}$ | — | < sliding cut † | — |
| p_T^{ratio} | — | $> 0.35^\dagger$ | — |

A few corrections are applied to all muons in simulation to address small differences between simulation and data. The first set of corrections are called Rochester corrections [167] and aim to bring the energy spectrum of muons in simulation closer to that observed in data. A second set of corrections comes in the form of scale factors on the efficiency of the identification. They correct the efficiency of the selection in simulation for that observed in data.

6.5.2 Electrons

In the same fashion as the muon ID, **loose electrons** should be found within the confines of the tracker, leading to a $|\eta|$ cut of 2.5, the impact parameters and their significance should be small (see Table 6.2) and the electron should be isolated. Additionally, the track can have no more than one missing hit. In order to avoid double counting and selecting electrons that are wrongly reconstructed from muon signatures, all electrons that are within a cone $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.05$ from loose muons are disregarded.

The **FO working point** requires a few additional selections based on the shape of the ECAL and HCAL clusters and the matching between the track and the ECAL cluster. For genuine electrons one expects most of the shower to be contained in the ECAL with only a small

punch-through to the HCAL. As such, we can look at the fraction of energy in the HCAL found within a cone of $\Delta R < 0.15$ around the ECAL cluster direction to the energy deposit in the ECAL, H/E . Secondly, a selection is made on the shower shape in the ECAL in the form of the $\sigma_{inj\eta}$ variable, the exact definition of which can be found in reference [168]. Additionally, the information between tracker and ECAL should agree. For this purpose, a cut on $1/E - 1/P$ compares the momentum measured at the closest point to the primary vertex P to the energy in the ECAL cluster E . One extra source of background electrons is found in photon conversions into two electrons. These are rejected using a conversion veto that requires there to be no tracks present in the tracker that are matched to a conversion vertex and point in the direction of the ECAL cluster.

Additional requirements are placed on the output score of an MVA trained in a DY MC sample using discriminating variables such as the ones described in the last paragraph. The purpose of this MVA is to provide a handle on discriminating prompt electrons from those misidentified from jets. It is different from the lepton MVA described in the muon section and the tight selection in the next paragraph and does not use isolation variables as inputs. It is labelled as the *electron MVA*. To top off the FO selection, a selection on the DEEPFLAVOR score of the closest jet and p_T^{ratio} is applied, the p_T threshold is raised to 10 GeV and there are no missing hits in the tracker allowed anymore.

The **tight working point** requires an electron to first pass the FO working point. On top of that, a selection is made on the same lepton MVA that was introduced in the muon section. Similarly as for the muons, energy scale corrections are applied to electrons in simulation along with scale factors on the identification efficiency. The first set will bring the reconstructed electron energy spectra in simulation and data closer to each other while the second takes slight differences in the behaviour of the identification selections between data and simulation into account. On top of this, reconstruction efficiency corrections are applied to match the efficiency of the electron reconstruction algorithms in simulation and data. The final selection efficiency for tight prompt electrons is 75% in the three-lepton selection with a misidentification rate of around 5% as measured for both with respect to the baseline selection.

6.5.3 Tau Leptons

As defined in Section 4.3.8, τ leptons are reconstructed from their hadronic decay products and will have similar characteristics to QCD jets. It is not uncommon that the HPS algorithm erroneously reconstructs a QCD jet as a τ_h jet. Therefore it is important to perform further quality selections on reconstructed τ candidates in order to retain as many of the genuine τ_h 's and remove as many of the misreconstructed τ_h 's as possible.

A baseline selection on the transverse momentum and pseudorapidity of τ_h 's of respectively $p_T > 20$ GeV and $|\eta| < 2.3$ performs a selection on the fiducial region for τ_h reconstruction. Furthermore, a decay mode finding discriminator selects only candidates with a vertex according to a 1-prong decay or a 3-prong that has no additional π^0 . In order to discriminate genuine

Table 6.2: Selection requirements that define the electron ID WPs. Cuts that are labelled with † are applied in the case that the electron does not meet the lepton MVA requirement.

| Cut | Loose | Fakeable Object (FO) | Tight |
|-----------------------------------|-------|--|---------|
| $ \eta < 2.5$ | ✓ | ✓ | ✓ |
| p_T (GeV) | > 7 | > 10 | > 10 |
| $ d_{xy} < 0.05$ cm | ✓ | ✓ | ✓ |
| $ d_z < 0.1$ cm | ✓ | ✓ | ✓ |
| $SIP_{3D} < 8$ | ✓ | ✓ | ✓ |
| $I_{mini}^{rel} < 0.4$ | ✓ | ✓ | ✓ |
| $H/E < 0.1$ | — | ✓ | ✓ |
| $\sigma_{in\eta} < (0.011/0.030)$ | — | ✓ | ✓ |
| $1/E - 1/P > -0.04$ | — | ✓ | ✓ |
| conversion veto | — | ✓ | ✓ |
| number of missing hits | < 2 | $= 0$ | $= 0$ |
| electron MVA | — | Loose (WP90 †) | Loose |
| lepton MVA | — | — | > 0.4 |
| DEEPFLAVOR ^{closest} | — | < tight WP | — |
| p_T^{ratio} | — | $> 0.6^\dagger(2018)/0.7^\dagger(2016 - 2017)$ | — |

τ 's against those whose signature was faked by a muon or an electron, all τ_h found within a cone $\Delta R = 0.4$ around loose light leptons are removed from the candidate collection.

At the time of the development of this analysis, a few different τ_h identification methods to be used on top of the baseline selection were available. Besides a cut-based identification method, a few iterations of MVAs were available. The DeepTau algorithm already referred to in Chapter 5 was still under development and considered experimental at the time. In order to probe which algorithm would give the best performance in light of the signal process under consideration, a comparison of efficiency versus misidentification rate was performed. The results of such a comparison in MC between a cut-based discriminator, the 2015 and 2017v2 versions of the MVA both for the old decay mode reconstruction and the new decay mode reconstruction (which contains 3-prong vertices with an additional π^0), and the deepTau discriminator can be found in Figure 6.2 for the definitions of efficiency and misidentification rate defined in the next paragraphs.

The efficiency is measured by first selecting all τ_h candidates flagged as genuine in simulation and that pass the baseline selection. One can then check how many of those pass the corresponding working point of the algorithm in question. The ratio between the number of candidates passing that working point and the number of τ 's passing the baseline selection is then taken as the efficiency. Because of the large range of models and signal masses, a WZ sample was used as a proxy for the signal in order to measure this efficiency. The definition of

the misidentification rate is defined in a similar fashion but instead of selecting genuine τ 's, all candidates must be not matched to τ leptons on a simulation truth level. The sample used for this measurement is a DY sample.

It can be seen from the receiver operating characteristic (ROC) curves in Figure 6.2 that the cut-based approach is outperformed by all other methods. Among the MVAs, a convergence in performance is observed when going to tighter working points. Lastly, the deepTau algorithm offers lower misidentification rate for the same efficiency but due to the experimental nature of this neural network at the time, it was decided to go with the 2017v2 MVA.

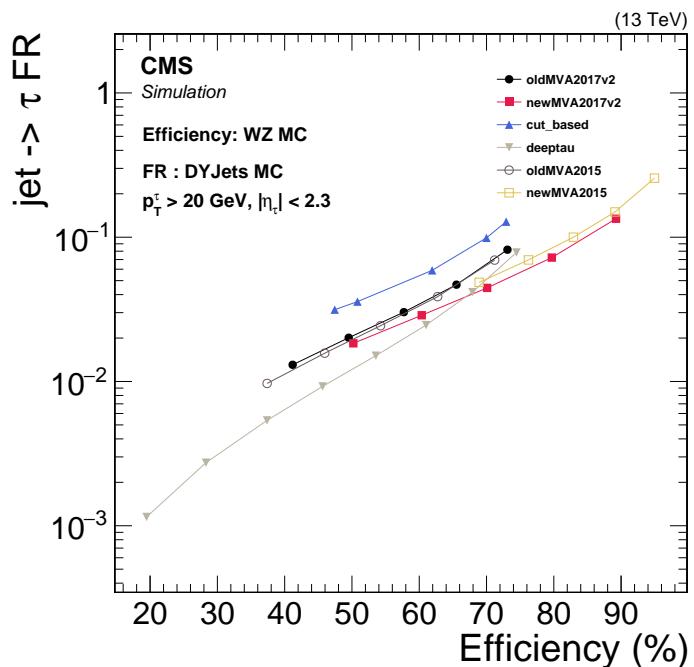


Figure 6.2: Efficiency and misidentification rate of several τ_h isolation ID algorithms as measured in 2016 Monte Carlo samples. Efficiency is measured in a WZ sample while misidentification rates are measured in a DY sample.

A study of the significance $S/\sqrt{S+B}$ where S is the signal yield and B is the total background yield as a function of all possible working point combinations was performed in order to choose which one offers the highest sensitivity to our signals. For this study, a variety of tau-enriched and tau-dominated signal samples was used. All events that go into the study need to pass the full event selection as listed in Section 6.6. The results of this comparison for a compressed scenario with $m_{\chi_1^0} = 150$ GeV and $m_{\chi_2^0} = 250$ GeV can be found in Figures 6.3. It can be seen that the tight working point of the MVA offers the best sensitivity balance for the single tau and ditau channels. The same study was performed with signal samples at different superpartner mass points. These studies showed that the tight working point offers the most balanced discrimination for both compressed and uncompressed scenarios and high sensitivity to final states with more than one τ_h . In order to remove most of the remaining τ_h that are faked by light leptons that passed the cleaning process from electrons and muons,

the loose working points of MVA discriminators against such leptons (called `againstElectron` and `againstMuon`) are applied. A summary of the full selection can be found in Table 6.3. The *loose* working point in this table signifies the working point to be used in nonprompt background estimation, which will be further explained in Section 6.7. The final selection, marked as *tight*, has a 50% efficiency in selecting prompt τ_h with a misidentification rate below 0.2% with respect to the baseline selection.

Table 6.3: Selection requirements that define the τ_h ID WPs.

| Cut | Loose τ_h selection | Tight τ_h selection |
|------------------------|--------------------------|--------------------------|
| $ \eta < 2.3$ | ✓ | ✓ |
| p_T (GeV) | > 20 | > 20 |
| againstMuon WP | Loose | Loose |
| againstElectron WP | Loose | Loose |
| MVA2017v2 isolation WP | Loose | Tight |

Energy scale corrections for genuine τ_h and τ_h faked by electrons that bring the energy profile of τ_h candidates in simulation closer to what is observed in data are applied to the relevant candidates in this analysis. Similarly, scale factors to correct for differences in the identification efficiency in data and simulation that are applied to genuine τ_h , muons misidentified as τ_h and electrons misidentified as τ_h .

6.5.4 Jets and b-tagging

All jets used in this analysis are PF jets reconstructed as explained in Section 4.3.6. Just like in the case of the objects discussed so far, a few identification requirements are applied to jets as well. Firstly, all jets must pass a set of standard jet identification requirements that are referred to as the Tight jet ID [169]. Furthermore, all jets where a FO electron, muon or τ_h was found within a cone of $\Delta R < 0.4$ centred on it are removed from the final jet collection. Additionally, the p_T of the jet should exceed 25 GeV and the candidate should be found within $|\eta| < 2.4$. All jets have the JEC and JER corrections mentioned in Section 4.3.6 applied.

Jets that originate from b-quarks are selected from the pool of identified jets and tagged using the DeepCSV algorithm [127]. The main purpose for b-tagged jets will be for vetoing events that contain such a jet. In order to correct for known differences in the b-tagging efficiency between simulation and data, scale factor weights are applied to every event in simulation.

6.5.5 Missing transverse energy

The missing momentum vector \vec{p}_T^{miss} is defined as the inverted vector sum of the transverse momentum of all PF candidates. The aforementioned JEC uncertainties are taken into account in this sum. The size of this vector is referred to as the missing transverse momentum p_T^{miss} . In order to remove anomalous high- p_T^{miss} events from reconstruction failures, detector malfunctions or non-collision backgrounds, MET filters are applied in this analysis.

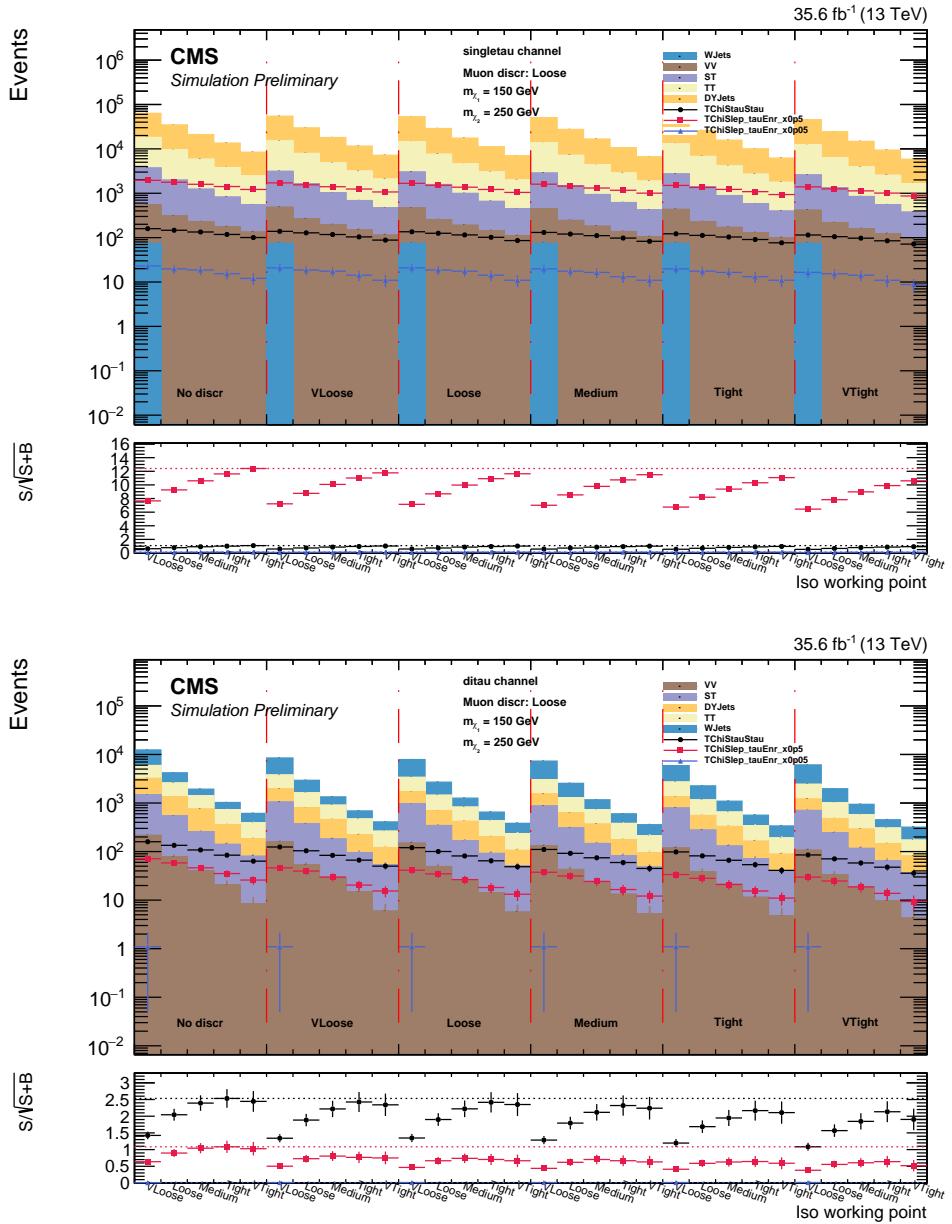


Figure 6.3: Number of $ll\tau_h$ events (top) and $l\tau_h\tau_h$ events (bottom) passing full event selection where the τ_h passed the according object selection in 2016 simulation. Different working points for the 2017v2 MVA are shown on the x-axis while different working points for the againstElectron discriminator are shown in different columns. In both cases, the loose working point of the againstMuon discriminator is applied. Background yields are represented in the filled-in distributions. The yields for three different slepton-mediated signal scenarios are shown as coloured points: τ -dominated in black, τ -enriched with $x = 0.5$ in red and τ -enriched with $x = 0.05$ in blue.

6.6 Analysis strategy and event selection

A multitude of final states is targeted in this analysis in order to be sensitive to a wide range of models. All of them share the presence of multiple leptons and a substantial amount of missing energy due to the LSPs.

The goal of the analysis is to use what we know about the signal properties to design variables that have different distributions for signal and background events, also called *discriminating variables*. By creating regions – collections of events that have to satisfy a number of well-defined selection requirements – from these discriminating variables that are highly enriched in expected signal while having low background yields, we can study the presence of a potential excess of events on top of the predicted background in the data. In case no such excess is observed, a fit to the data is performed in order to produce exclusion limits on the signal cross section in the relevant parameter space.

As a baseline selection, all events should pass at least one of the single light lepton or double light lepton triggers used at the time of data-taking. These triggers each come with their own online thresholds that define the turn-on curve. In order to be coherent in the offline and online selection, a number of p_T cuts are applied that make sure we reside on the plateau of constant efficiency. Additionally, events must have at least three FO leptons or two FO light leptons of the same sign. Events where these leptons pass the tight working point are kept to be used in the signal regions as the nominal selection. Events where one of the leptons fails the tight requirement are retained for further use in the data-driven background estimation methods. In order to reduce the backgrounds that contain top quarks, a veto on the presence of a b-tagged jet is put in place. Similarly, events that contain an opposite sign and same flavour (OSSF) pair that has a combined invariant mass that does not exceed 12 GeV are discarded to decrease contributions from photon conversions and low-mass resonances.

All surviving events eventually get distributed into different *search categories* based on the number of leptons, flavour content and the number of OSSF lepton pairs, for each of which a set of dedicated search regions are designed. As mentioned in the first sections, we will limit the scope of this chapter to the final states relevant to the τ_h studies performed by the author, namely trilepton events including hadronically decayed τ 's. However, it should be acknowledged that there are other final states considered in the analysis, such as two same-sign light leptons or three light leptons. In fact a number of interesting developments, such as a novel parametric neural network technique, have been made for these final states and they offer the bulk of the sensitivity to the flavour-democratic scenario of slepton-mediated decay models. Be that as it may, these final states and their results will not be discussed any further. The interested reader is referred to the paper instead [7].

Three charged leptons with one or more tau leptons

Channels with three charged leptons of which at least one is a τ_h are especially sensitive to the tau-enriched and tau-dominated slepton-mediated models. While it is possible that τ leptons decay leptonically and thus help form a final state with only light leptons, the bulk of the sensitivity for these two types of models will come from final states with at least one τ_h . The exact final states in this category used in this analysis can be split into three subcategories:

- One OSSF light lepton pair + τ_h ($3\ell C$):

These events are very sensitive to τ -enriched models but come at the price of high

background yields. As will be shown in Section 6.7.1, DY events with an additional jet can fake this signature. This background can already largely be reduced by vetoing events where the light lepton pair has an invariant mass within 15 GeV of the Z boson mass, however. A second large background comes from $t\bar{t}$ events with a nonprompt τ_h .

For events with p_T^{miss} below 300 GeV, this background is reduced with cuts on a variable referred to as the stransverse mass, M_{T2} [170]. It is a complex variable designed for processes with two semi-invisible decays of mother particles with the same mass. In order to understand it, we need to first define the transverse mass variable M_T , which is designed for processes with a single semi-invisible decay. For a system with a lepton and missing transverse momentum, it is defined as follows:

$$M_T = \sqrt{2p_T^\ell p_T^{\text{miss}}(1 - \cos(\Delta\phi))}, \quad (6.4)$$

where p_T^ℓ is the transverse momentum of the lepton and $\Delta\phi$ is the angle between the lepton and the missing momentum vector in the transverse plane. This quantity will by definition have a smaller value than the mass of the mother particle of the semi-invisible decay and will thus have a sharp drop-off at the mass of the mother particle, making it effective against backgrounds that contain for example leptonic W decay.

In a similar fashion, the stransverse mass is calculated in such a way that it falls off sharply beyond the mother particle mass in processes with two semi-invisible decays. In short, this is done by running over all possible splittings of the total missing transverse momentum vector and minimizing the maximum value of each decay's transverse mass. For more details and mathematical background, the reader is referred to the reference material.

For higher missing energy values we depend on the transverse mass of the combined system of the two light leptons and p_T^{miss} (M_T^{2l}) for this discrimination. The exact definitions of the search regions can be found in Table 6.4.

- Non-OSSF light lepton + τ_h ($3\ell D$ - $3\ell E$):

This subcategory is in itself split again into two categories, namely events where the light lepton pair have opposite sign (OS) and those where they do not. Events with opposite sign and different flavour (OSDF) light leptons are categorized under category $3\ell D$ while events with two same sign light leptons fall under $3\ell E$. In scenarios where a Z boson decays to two τ 's, the event can contribute to these final states if either one or both τ 's decay leptonically. In that case, the reconstructed mass would not be the Z boson mass because the neutrinos in the τ decay escape with a part of the momentum. It has been shown that the expected reconstructed mass for such Z decay is around 50 GeV for $e\mu$ final states and around 60 GeV for $1\tau_h$ final states. With this in mind, search region bins as a function of the invariant mass of the lepton pair with opposite sign are created. In the case of OS light leptons, this will be their invariant mass ($M_{\ell\ell}$).

In case the light leptons have the same sign, this will be the invariant mass of the light lepton and τ_h combination that is closest to the expected value of 60 GeV ($M_{\ell\tau_h}$). When the light leptons have the same sign and also share the same sign with the τ_h candidate, this mass variable is set to 0 GeV. Lastly, the stransverse mass M_{T2} is added into the mix. It is calculated using the two light leptons if they have OS or the leading light lepton³ and the τ_h in the other case. Due to the sharp drop in the distribution beyond the W mass, it adds a lot of discriminating power to this category. The exact definitions of the search regions can be found in Tables 6.5 and 6.6.

- Two τ_h candidates (3 ℓ F):

Events that contain two τ_h candidates are especially sensitive to the tau-dominated slepton mediated decay models. The two variables that are the foundation of the binning for this category are based on the leading τ_h and the light lepton. The first of these is the invariant mass of this pair $M_{\ell\tau_h}$, which tends to be large for uncompressed scenarios while the second is the stransverse mass of the pair in combination with the missing transverse energy vector, $M_{T2}(\ell\tau_h)$. The exact definitions of the search regions can be found in Table 6.7.

Table 6.4: Definition of the search regions for events with an OSSF light lepton pair and a τ_h (3 ℓ C).

| p_T^{miss} (GeV) | $M_T^{2\ell}$ (GeV) | $M_{T2} < 80$ GeV | $80 \leq M_{T2} < 120$ GeV | $M_{T2} \geq 120$ GeV |
|---------------------------|---------------------|-------------------|----------------------------|-----------------------|
| 50–200 | ≥ 0 | C01 | C02 | C03 |
| 200–300 | ≥ 0 | C04 | C05 | C06 |
| | 0–250 | | C07 | |
| ≥ 300 | 250–500 | | C08 | |
| | ≥ 500 | | C09 | |

Table 6.5: Definition of the search regions for events with an OSDF light lepton pair and a τ_h (3 ℓ D).

| $M_{T2}(\ell, \ell)$ (GeV) | p_T^{miss} (GeV) | $M_{\ell\ell} < 60$ GeV | $60 \leq M_{\ell\ell} < 100$ GeV | $M_{\ell\ell} \geq 100$ GeV |
|----------------------------|---------------------------|-------------------------|----------------------------------|-----------------------------|
| 0–100 | 50–100 | D01 | D06 | D11 |
| | 100–150 | D02 | D07 | D12 |
| | 150–200 | D03 | D08 | D13 |
| | 200–250 | D04 | D09 | D14 |
| | ≥ 250 | D05 | D10 | |
| ≥ 100 | 50–200 | | D15 | |
| | ≥ 200 | | D16 | |

³The term leading lepton refers to the lepton with the highest p_T while subleading and trailing leptons respectively refer to the lepton with the second highest and third highest p_T .

Table 6.6: Definition of the search regions for events with a same-sign light lepton pair and a τ_h ($3\ell E$).

| $M_{T2}(\ell, \tau_h)$ (GeV) | p_T^{miss} (GeV) | $M_{\ell\tau_h} \leq 50$ GeV | $50 < M_{\ell\tau_h} \leq 100$ GeV | $M_{\ell\tau_h} > 100$ GeV |
|------------------------------|---------------------------|------------------------------|------------------------------------|----------------------------|
| 0–80 | 50–100 | E01 | | E04 |
| | 100–250 | E02 | | E05 |
| | ≥ 250 | E03 | | |
| ≥ 80 | 50–150 | | E06 | |
| | 150–200 | | E07 | E08 |
| | ≥ 200 | | | E09 |

Table 6.7: Definition of the search regions for events with two τ_h candidates and a light lepton ($3\ell F$).

| $M_{T2}(\ell, \tau_h)$ (GeV) | p_T^{miss} (GeV) | $M_{\ell\tau_h} < 100$ GeV | $M_{\ell\tau_h} \geq 100$ GeV |
|------------------------------|---------------------------|----------------------------|-------------------------------|
| 0–100 | 50–100 | F01 | F07 |
| | 100–250 | F02 | F08 |
| | 150–200 | F03 | F09 |
| | 200–250 | F04 | F10 |
| | 250–300 | F05 | |
| | ≥ 300 | F06 | |
| ≥ 100 | 50–200 | | F11 |
| | ≥ 200 | | F12 |

6.7 Background Estimation

Even after all the effort of reducing the background contributions to the final yields, there will always be some level of background creeping through. It is now of utmost importance to understand what these backgrounds are and to accurately predict them. In this search specifically, there will be four main types of background contributions to our final selection. The first of these are events where one of the selected leptons is actually nonprompt. An example of this would be a DY event that decays to two OSSF light leptons but in which a final state radiation jet mimics a hadronically decayed tau. While there are no three prompt leptons, when the nonprompt lepton slips through our selection, it will look like there are. The background due to nonprompt leptons forms the main background for the trilepton channels that include a τ_h . To predict it, a data-driven technique is employed. We will go a bit more into detail on this technique in the following subsection.

A second dominant type of background is due to SM processes with the same final state, such as a WZ or ZZ process ending up in 3 or 4 lepton final states. The contribution from prompt SM processes is predicted from simulation. Conversions make up the third type of background. They are events in which a photon evolves into an electron-positron pair. Two mechanisms through which this can happen exist: internal and external conversion. The first is simply the splitting of a virtual photon while the second is due to a real photon interacting with the detector material. Just like the prompt background processes, we will predict its contribution from simulation.

6.7.1 Nonprompt background processes

Data-driven background estimation methods are employed for the prediction of the contribution of this type of background. More specifically, a fake rate or “tight-to-loose” method forms the foundation of the prediction in this background category. While the specific measurement and application will differ for nonprompt light leptons and nonprompt τ_h , the main principle is the same. One can select a control region in the data that is orthogonal to the signal region selection and is enriched in nonprompt leptons of the flavour in question. By counting the number of nonprompt leptons that pass the FO selection and the number of those leptons that then also pass the tight analysis object selection, a map can be formed of the probability for a nonprompt lepton to pass the full selection. This map is called the *fake rate* and is for our intents and purposes a function of the lepton p_T and $|\eta|$.

With this fake rate map in hand, one can estimate the contamination of nonprompt backgrounds in the signal region by selecting a sideband region in data. The latter is defined similarly to the signal region with the exception that at least one of the leptons should pass the FO selection and not the tight selection. One can then define and apply a transfer factor that is a function of the fake rate to translate the events in the sideband to the contribution of nonprompt backgrounds to the signal region.

In the hypothetical scenario where there can only be one nonprompt lepton and therefore just one lepton in the sideband that fails the tight selection but passes the FO selection, this transfer factor is found to be

$$w = \frac{f(p_T, |\eta|)}{1 - f(p_T, |\eta|)} , \quad (6.5)$$

where $f(p_T, |\eta|)$ represents the fake rate of a lepton with transverse momentum p_T and pseudorapidity η . Note that the denominator in the transfer factor is needed because of the definition of the denominator of the fake rate, namely all FO nonprompt leptons regardless of their passing of the tight working point. As the sideband is defined as consisting of FO leptons that do not pass the tight working point, the denominator in equation 6.5 essentially translates the “tight-to-loose” ratio into a “tight-to-loose-but-not-tight” ratio. This can be shown as follows: if the number of nonprompt leptons passing the tight selection is denoted by $N(t)$, the number of nonprompt leptons passing the loose selection by $N(l)$ and the number of nonprompt leptons passing the loose selection but not the tight selection by $N(l|!t)$, the transfer factor can be written down as:

$$w = \frac{N(t)}{N(l|!t)} = \frac{N(t)}{N(l) - N(t)} = \frac{\frac{N(t)}{N(l)}}{1 - \frac{N(t)}{N(l)}} = \frac{f}{1 - f} . \quad (6.6)$$

In the scenario where it is possible that there are multiple nonprompt leptons, the transfer factor to apply as a weight to the event becomes more complicated since one must take into account additional corrections due to combinatorics [171]. Imagine two-lepton events where the signal region of two tight leptons can have background contributions from events with one or two nonprompt leptons. The first can be estimated from the sideband of events with one

tight lepton and one FO-but-not-tight lepton by applying the transfer factor $f_l/(1 - f_l)$ we have already encountered. The “problem” is that this sideband selection is not pure in events with one nonprompt lepton but also contains events with two nonprompt leptons where one of them actually passed the tight selection. At first sight, this is a good thing because it therefore provides a prediction of both the contribution from one-nonprompt and two-nonprompt lepton events. However, for the latter type, there are two combinations in which one lepton passes the tight working point and thus there is a double counting of events with two fake leptons.

This double counting can be counteracted by estimating the contribution of events with two nonprompt leptons from a sideband of events with two FO-but-not-tight leptons (where there is no combinatorial ambiguity) and subtracting it from our estimation from the sideband with one FO-but-not-tight lepton and a tight lepton. The final weight to be applied to the event can be factorised and it can be shown that it can be generalized for multilepton final states as follows [172]:

$$w = (-1) \times \prod_l \frac{-f_l}{1 - f_l}, \quad (6.7)$$

where the product runs over all leptons that fail the tight selection but pass the loose selection and f_l is the fake rate at the p_T and $|\eta|$ of the according lepton.

Nonprompt light leptons

In order to measure the fake rate maps for nonprompt light leptons, a QCD multijet region is selected in data that is enriched in this type of lepton. Before measuring the fake rate, there are first two effects that need to be taken into account. Firstly there is the matter of the fake rate dependency on the p_T of the mother parton that was the source of the jet that contains the nonprompt light lepton. The transverse momentum that is reconstructed and assigned to the nonprompt lepton is not the same as the p_T of that mother parton. If the light lepton is, for example, a nonprompt lepton that is contained in a QCD jet, it will only contain a fraction of the energy of the original parton. Imagine that for the same reconstructed p_T , one lepton originates from a low p_T parton while another originates from a high p_T parton. The activity around the second lepton is likely to be higher than in the case of the first lepton. This will of course have an impact on the probability of passing the full selection for which we highly depend on isolation variables. Secondly, there is the matter of the nature of the mother parton. Nonprompt light leptons in the signal region come mainly from DY and $t\bar{t}$ processes that respectively have more fakes from light flavour quarks and heavy flavour quarks. The different structures in their resulting jets will cause different behaviour in their fake rates. In order to take this into account, a working point called the fakeable object working point is defined as providing a similar tight-to-FO ratio for heavy and light flavour jets.

To address the first point, a so-called “cone-correction” that translates the nonprompt lepton p_T to a proxy of the mother parton p_T is applied to all nonprompt leptons, i.e. leptons that pass the FO selection but fail the tight selection. A good proxy for the mother parton p_T is

the p_T of the resulting jet, which can be obtained by dividing the original p_T with the earlier defined p_T^{ratio} . That means that this cone-correction comes down to multiplying the nonprompt lepton p_T by a factor $\frac{x}{p_T^{\text{ratio}}}$ where x is a scale factor that corrects for the discontinuity in the resulting transverse momentum spectrum of the leptons. This discontinuity can be attributed to the fact that the correction is only applied to leptons that fail the tight selection. If one was to plot the average value of the cone-corrected p_T (p_T^{cone}) as a function of the lepton MVA used for light lepton identification, one would see a jump in the spectrum as demonstrated in Figure 6.4. Therefore this value x can be obtained by rescaling the distribution below the tight selection cut to smoothly connect to the rest of the distribution again. For muons and electrons, these values were found to be 0.65 and 0.85 respectively.

The second point is addressed by defining a large range of FO definitions and probing which one provides hadron-independent fake rates. In practice, this is done by measuring the tight-to-loose ratios for each of those FO definitions in a simulated QCD sample and checking how well they predict the number of events in both $t\bar{t}$ and DY samples that have three tight leptons from a sideband of events where at least one light lepton fails the tight working point (but passes the FO working point). The definition of the FO working point that provided the best hadron-independent predictive power was already defined in Tables 6.1 and 6.2.

After performing these measurements in simulation as described in the previous paragraph, the tight-to-loose ratio maps were measured in the same QCD multi-jet region in data. The contribution in data due to prompt light leptons was estimated from simulation and subtracted from both the numerator and the denominator. The resulting tight-to-loose ratio maps can be found in Figure 6.5 as a function of p_T and $|\eta|$. Due to the high contamination of prompt leptons above a p_T value of 45 GeV, the tight-to-loose maps are truncated at this value. Comparison with the fake rate maps obtained from simulation shows good agreement for the bins below this truncation. It also shows that for values above this truncation, the fake rate is not completely flat. However, the contribution from nonprompt light leptons with a p_T above 45 GeV is relatively small and closure tests show that this effect is covered by a 30%

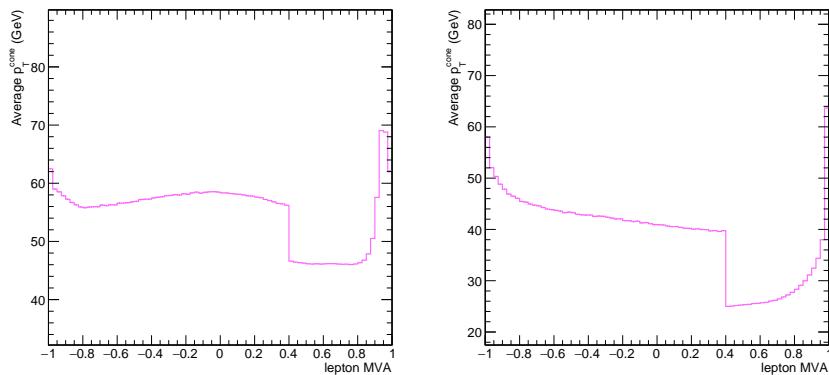


Figure 6.4: The average value of p_T^{cone} without additional correction factor as a function of the leptonMVA for electrons (left) and muons (right) as measured in QCD simulation. Plots by Willem Verbeke.

systematic uncertainty.

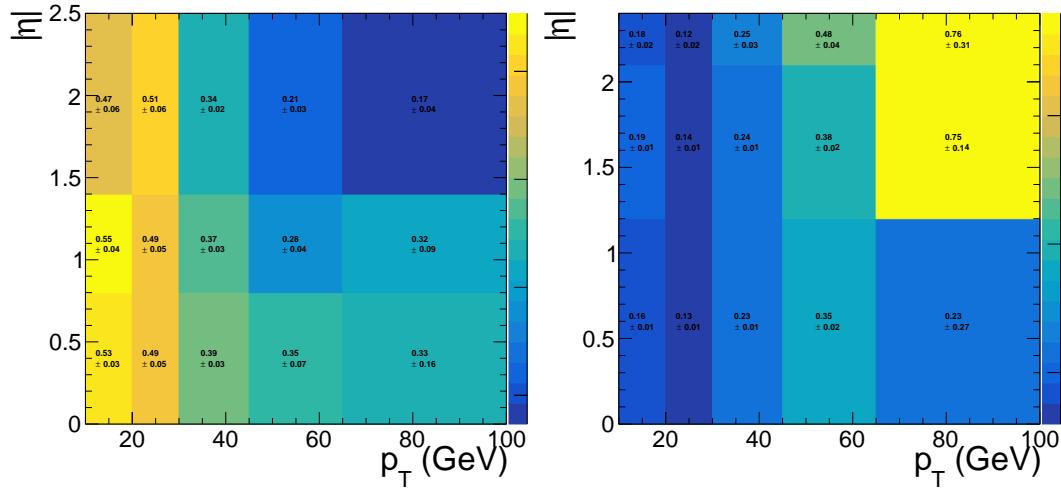


Figure 6.5: Tight-to-loose ratios for nonprompt electron (left) and muon (right) estimation as measured in a QCD control region using data corresponding to 2017 data-taking. Plots by Willem Verbeke.

Nonprompt taus

While the fake rate for τ_h 's is also dependent on the origin of the object faking the τ_h , it is not possible to solve this issue in the same way as for nonprompt light leptons. The loose working point used for τ_h identification is completely bound by the definition of MVA working points and is not easily tweaked by the user to have similar fake rates for light jet and heavy jet fakes. To facilitate this difference in fake rate for each origin, which mainly comes from the different flavour content in these two types of jets, not one but two fake rate measurement regions were defined. One of these is a DY enriched region where the main source of τ_h fakes comes from light flavour jets. The second control region selects mainly on $t\bar{t}$ events and their corresponding heavy jet activity.

The DY control region is characterized by a Z boson in the form of the presence of two OSSF light leptons passing the full object selection. The invariant mass of this pair is required to have a value within 15 GeV of the mass of the Z boson. In addition to this light lepton pair, a τ_h candidate that passes the loose selection as defined in Table 6.3 should be present in the event. Outside of these three lepton candidates, there can be no other loose τ_h or FO light lepton in the event for it to be retained in this control region. Finally, the missing transverse momentum is required to be less than 50 GeV in order to remove contributions with a prompt τ_h .

The $t\bar{t}$ control region is formed by selecting events with a single electron, a single muon and a single τ_h . Similarly to the DY control region, the light leptons are required to pass the tight selection while the τ_h should pass the loose selection. Any additional FO leptons imply a veto of the event. Contrary to the DY region, this region is required to have at least 1 b-tagged jet passing the tight working point of the DeepCSV algorithm. Furthermore, the invariant mass

of the electron-muon pair must be larger than 20 GeV to suppress QCD events.

As a proof of concept, the measurement is first performed and validated in simulated samples in order to verify the efficacy of the method in the context of nonprompt τ_h estimation. To remove the prompt contribution in these measurements, all selected τ_h are required to be nonprompt on a simulation truth level. The fake rates obtained in both control regions in simulation can be found in Figure 6.6 where it can be seen that there is indeed a small difference between the heavy flavour jets and light flavour jets fake rate with lower fake rate values in the $t\bar{t}$ region. Similar values of the fake rate were found in the other years of data-taking.

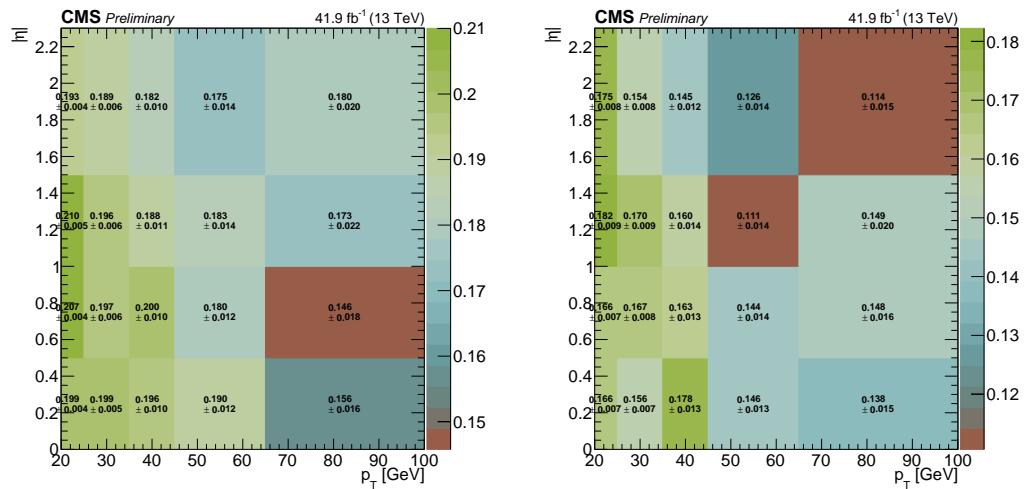


Figure 6.6: Tight-to-loose ratios for hadronic τ from a DY (left) and $t\bar{t}$ control region (right) dominated by fake τ_h from jets. Given for simulation samples corresponding to 2017 data-taking.

These measurements are subsequently tested with the use of closure tests in a τ_h -fake enriched control region. This means that a sideband region of events with at least one τ_h that fails the tight working point is selected and used to predict the number of events where all τ_h 's pass the tight working point. The light flavour jet closure tests in a DY region were performed in a similar region as the fake rate measurement region but with an inverted p_T^{miss} cut while the $t\bar{t}$ control region remains the same. Additionally, there should always be at least one τ_h in the event that is nonprompt on a simulation truth level. The results of these tests as performed in simulation can be found in Figures 6.7 and 6.8 for the categories relevant to the search.

After observing the efficacy of the method for τ_h 's in simulation, the tight-to-loose map measurements and subsequent closure tests are performed in data. The control regions used for this have the same definitions as the measurements in simulation. In order to remove contamination from prompt τ_h 's, their contribution is measured in simulation and subtracted from the data. In all closure tests the prompt contribution is also plotted on top of the “cleaned” prediction. The tight-to-loose ratio maps in 2017 data can be found in Figure 6.9 and their closure tests in 2017 data in Figures 6.10 and 6.11. The same behaviour was observed in 2016 and 2018 data.

In the final strategy, the choice of which fake rate map to use for nonprompt background estimation is different for each search region and is made based on the expected dominating nonprompt τ_h source in the region. In search regions 3lD and 3lE, $t\bar{t}$ processes are expected to make the largest contributions and thus the $t\bar{t}$ fake rate maps are used for the final background estimation. The story is different in search regions 3lC and 3lF, where we expect DY to be the leading contributor. As such the DY fake rate maps are used on those search regions.

6.7.2 Prompt backgrounds

When studying prompt background sources in the 3 lepton channels, it becomes clear that the dominant contributions come from WZ and ZZ processes respectively where the W decays to a charged lepton and a neutrino while the Z boson decays to a lepton-antilepton pair. Both of these are predicted from simulation in dedicated control regions designed to be depleted from signal process events. For the WZ process, this region is a subset of the search regions with three light leptons containing an OSSF pair. On top of the existing selection criteria, the following requirements are made:

- $|M_{\ell\ell} - M_Z| < 15 \text{ GeV}$
- $50 \text{ GeV} < p_T^{\text{miss}} < 100 \text{ GeV}$
- $50 \text{ GeV} < M_T^W < 100 \text{ GeV}$
- $|M_{3\ell} - M_Z| > 15 \text{ GeV}$

In order to get a normalization factor that can be applied to the WZ simulation, a fit is performed to the data in this region with all relevant uncertainties taken into account and in the assumption there is no signal present. From this fit a factor of 1.17 ± 0.05 is obtained.

In similar fashion to the WZ case, a control region was designed to select ZZ events and determine a normalisation factor. Again a fit is performed that includes all uncertainties and assumes there to be no signal while the ZZ normalisation is allowed to float freely. The normalisation factor obtained from this fit is 1.02 ± 0.10 .

Further prompt background contributions come from processes with one or more top quarks and electroweak bosons where the largest contributions come from events with a top quark-antiquark pair and a boson ($t\bar{t}Z$, $t\bar{t}W$, $t\bar{t}h$) along with the double top-antitop quark ($t\bar{t}t\bar{t}$) – also called four top – production. Collectively these will be labelled as $t\bar{t}X$. Processes with a single top quark and one or more electroweak bosons, labelled tX , contribute in lower amounts due to their small cross sections. Even smaller contributions come from events with a top quark-antiquark pair and two electroweak bosons, labelled $t\bar{t}XX$. All of these are very suppressed by the b-tagged jet veto in the analysis and are estimated from simulation.

We can also expect prompt contributions from events with three or more electroweak bosons. However, these processes are rare and only contribute in very small quantities. They are estimated from simulation.

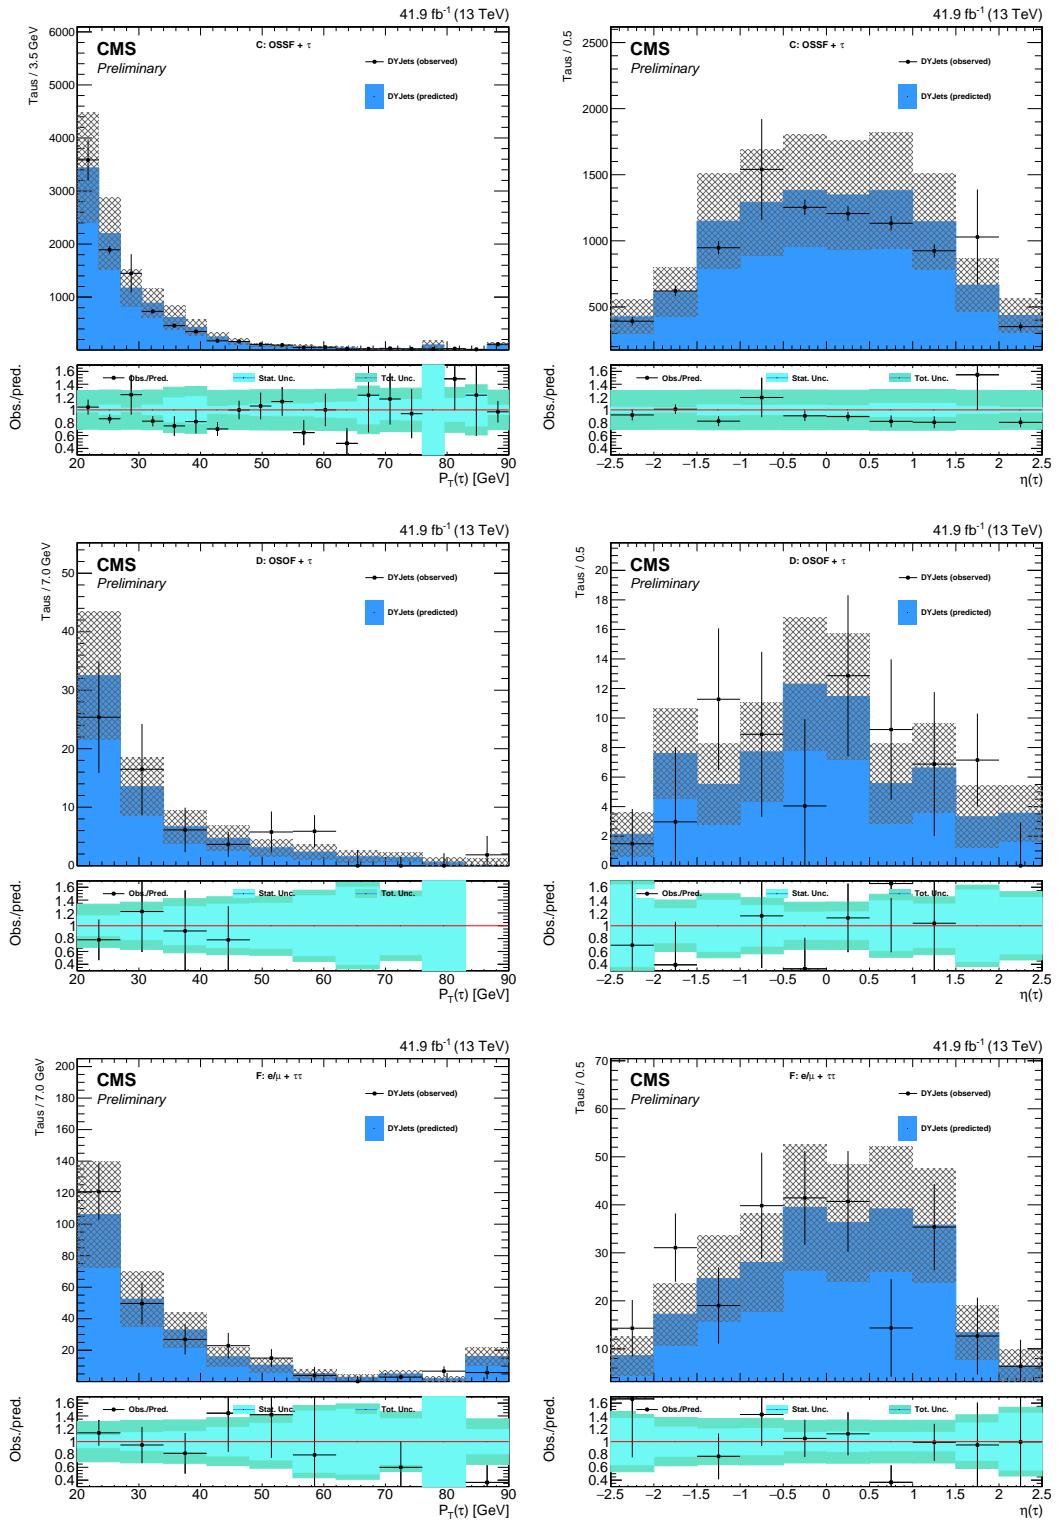


Figure 6.7: Closure tests performed in a 2017 DY MC sample using the fake rate measured in the DY control region for category 3l-C ($1^+1^-\tau_h$) (top), category 3l-D ($e\mu\tau_h$) (middle) and category 3l-F ($1\tau_h\tau_h$) (bottom) events as a function of $\tau_h p_T$ (left) and η (right) of all selected τ_h 's.

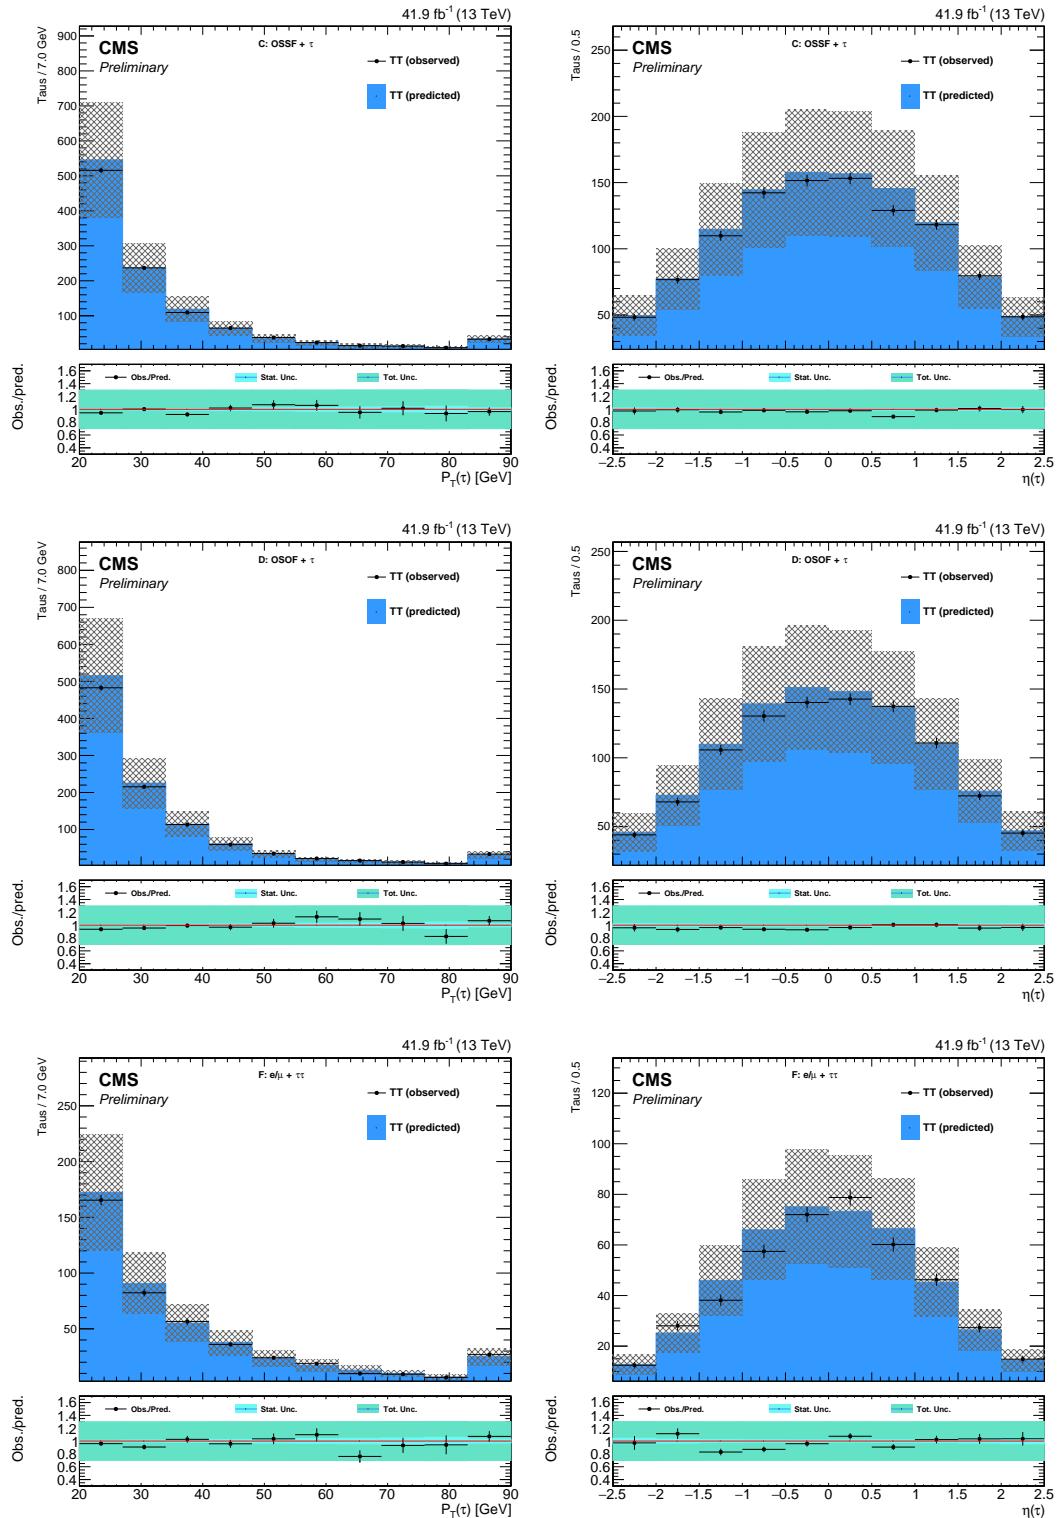


Figure 6.8: Closure tests performed in a 2017 $t\bar{t}$ MC sample using the fake rate measured in the $t\bar{t}$ control region for category 3l-C ($l^+l^-\tau_h$) (top), category 3l-D ($e\mu\tau_h$) (middle) and category 3l-F ($l\tau_h\tau_h$) (bottom) events as a function of τ_h p_T (left) and η (right) of all selected τ_h 's.

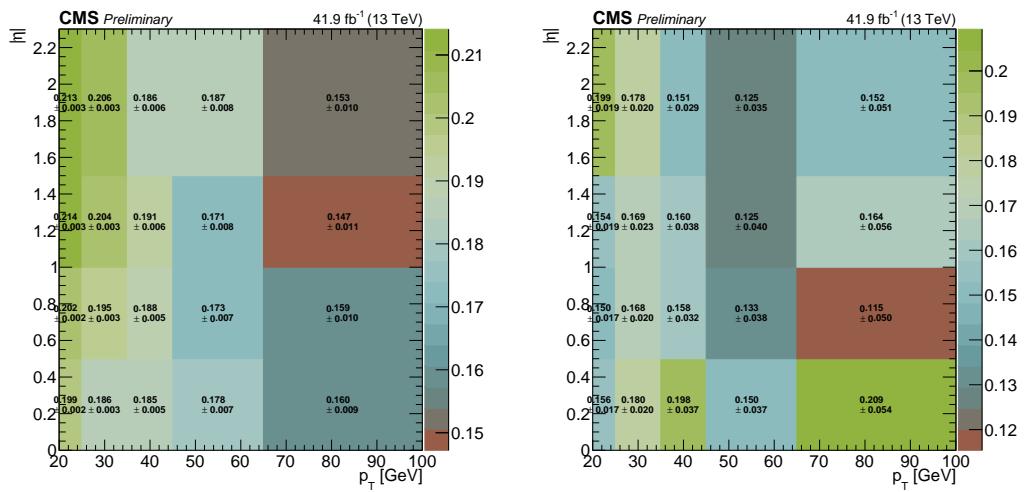


Figure 6.9: Tight-to-loose ratios for hadronic τ from a DY (left) and $t\bar{t}$ control region (right) dominated by fake τ_h from jets. Given for 2017 data.

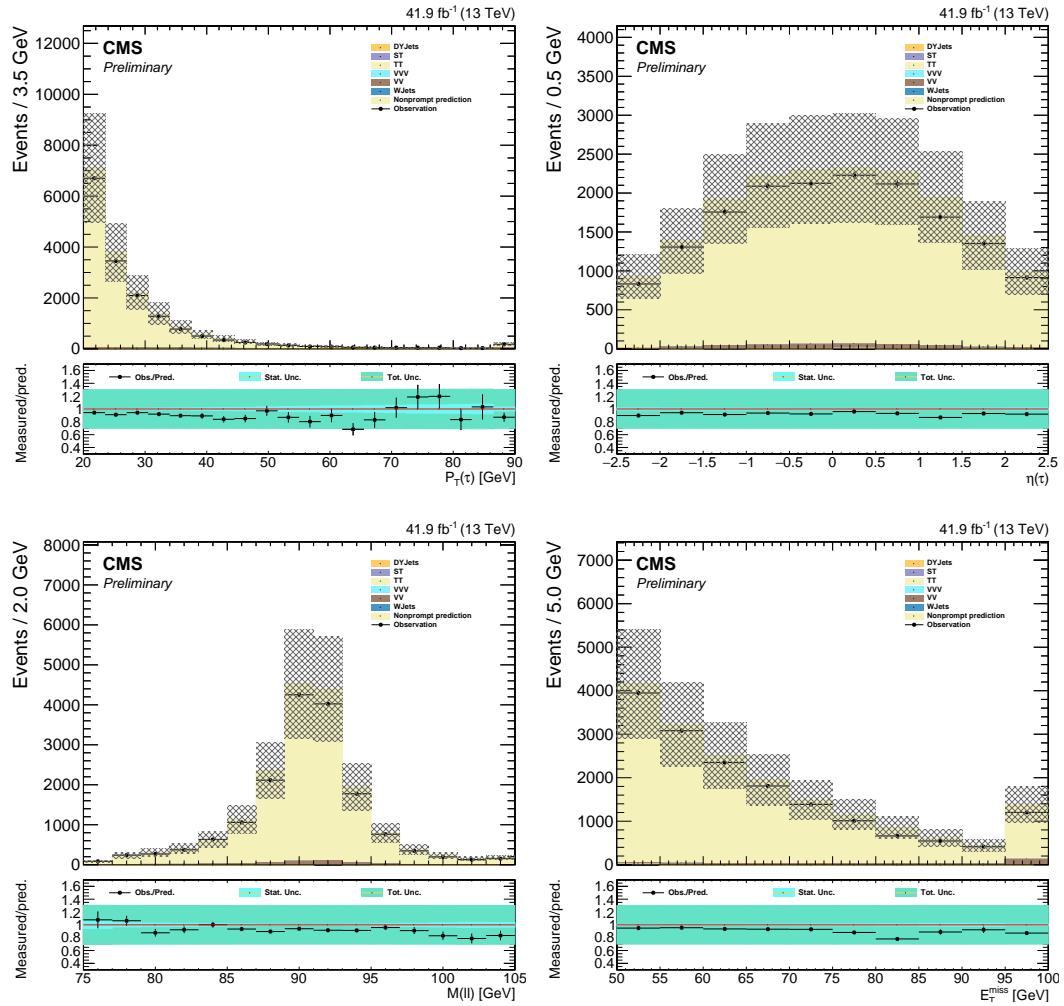


Figure 6.10: Closure tests performed in 2017 data in a control region with an OSSF light lepton pair with an invariant mass within a 15 GeV mass window to the Z boson mass and a single hadronic tau. A veto on b-tagged jets is applied as well as a cut of $E_T^{\text{miss}} > 50$ GeV. The contribution of prompt τ 's to the observation has been estimated from MC. Shown as a function of $\tau_h p_T$ (top left), $\tau_h \eta$ (top right), mass of the light lepton pair $M_{\ell\ell}$ (bottom left) and E_T^{miss} (bottom right).

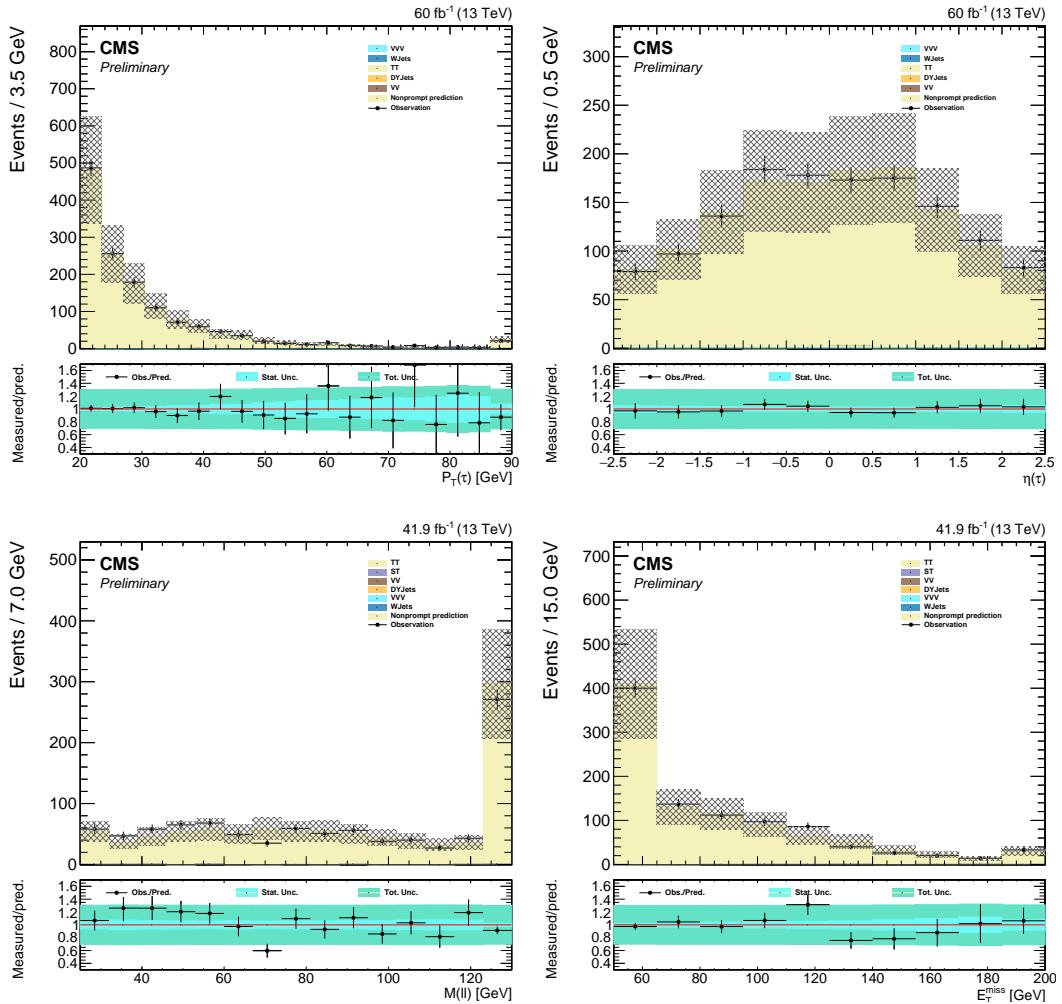


Figure 6.11: Closure tests performed in 2017 data in a control region with an electron, a muon and a hadronic tau. The two light leptons are required to have OS and to have a combined invariant mass above 20 GeV. At least 1 b-tagged jet needs to be present in the event. The contribution of prompt τ 's to the observation has been estimated from MC. Shown as a function of $\tau_h p_T$ (top left), $\tau_h \eta$ (top right), mass of the light lepton pair $M_{\ell\ell}$ (bottom left) and p_T^{miss} (bottom right)

6.7.3 Photon conversions

Processes where a photon converts either internally or externally to two leptons creep into all analysis categories. For the categories with three or more leptons the main contribution comes from $Z\gamma$ events where the Z boson decays leptonically. When the conversion happens asymmetrically in such a way that one lepton gets most of the energy while the other has very low p_T , it is possible that this latter lepton never gets reconstructed or identified, causing the event to be a trilepton event. This background type is estimated from simulation and validated in a dedicated control region that selects on events with low missing energy, a trilepton mass within 15 GeV of the Z boson mass and no OSSF pair with an invariant mass within 15 GeV of the Z boson mass. A normalisation factor of 1.12 ± 0.10 is obtained, which is applied to all conversion background simulation samples.

6.8 Systematic uncertainties

There are two types of uncertainty in the analysis. The first is a random uncertainty due to the statistical fluctuations that arise from the limited number of events in the simulation samples and the data. Its randomness comes from the fact that if one was to perform the same measurement multiple times, the yields will vary around the true value in a random way. The second type of uncertainty is labelled as systematic uncertainty. They are not statistical in nature but arise from limitations in the experiment. All aspects of performing an analysis, such as gathering data, running simulations, and predicting the expected background, inherently come with systematic uncertainties. A certain background process simulation could for example come with an uncertainty on the cross section and the true cross section will lie in a range with a size characterized by the size of the systematic uncertainty around this input cross section. This means that we are consistently slightly over- or underpredicting events belonging to that background. It is of utmost importance to locate all systematic uncertainties in the analysis and estimate their size and impact on the yields. These uncertainties can then be provided to our final fit where the potential variations are taken into account.

This analysis is affected by systematics regarding the simulated samples, systematics regarding uncertainties in the theoretical calculations that form the foundation of the simulation of the signal samples and systematic uncertainties on the data-driven background estimation. For each source of systematic uncertainty, the predicted signal and background yields are varied up and down in the relevant distributions and these variations are used to construct nuisance parameters in the fit. Below, all systematic uncertainties relevant to this analysis are listed and briefly discussed. You will see that there are two main categories of uncertainties described, namely *flat uncertainties* and *shape uncertainties*. With a flat uncertainty, we mean that the uncertainty only affects the total yield or equivalently that every bin in the search region bins will be affected in the same way. When a flat uncertainty number is quoted, it means that all bins are varied up and down with that factor. Shape uncertainties on the other hand affect each bin differently. Say the event gets weighted with a certain scale factor in the analysis, that weight for each event is scaled up and down to get the uncertainty. As this scaling will

potentially be different for different events, there might be trends in the uncertainties of the input shapes. Another concept that pops up is *correlation* between years. This simply refers to whether the systematic is independent for each year separately or not. In the final fit, the distributions of all three years form separate inputs to the simultaneous fit. If a systematic is independent, e.g. when it comes from a statistical limitation of the method which will be independent for each year depending on the available statistics in your samples, the systematic is called uncorrelated and is treated as a separate nuisance per year that is disabled for the other years in the fit. If a systematic uncertainty is correlated, the source of the systematic is shared between the years, e.g. stems from a method that is equivalently used across the years, the systematic is treated as a single nuisance that is turned on for all years at the same time. The most important systematic uncertainties are the following.

- The uncertainty on the **luminosity** measurement used to scale simulation sample yields to those expected during data-taking is a shared uncertainty among all CMS analyses. It is a flat uncertainty that is different for each year of data-taking and is partially correlated among different years. The size of the luminosity uncertainties is of the order of 2.5%.
- The **pileup** distribution as a function of the number of true interactions per bunch crossing during data-taking depends on the instantaneous luminosity, which is not always known when producing the simulated samples. Using the measured luminosity and a total inelastic proton-proton interaction cross section of 69.2 mb, the expected pileup distribution is estimated in data [163]. Simulation events can then be reweighted so that this distribution in the sample matches the expected one. To get weights for the up and down variations, the inelastic cross section used for the estimation of the expected distribution is scaled up and down by 4.6%.
- In order to gauge the agreement in **trigger efficiency** between data and MC, the efficiency of the combined light lepton triggers used in the analysis was measured in an unbiased, orthogonal event set that triggers on p_T^{miss} and hadronic momentum. The uncertainty on these measurements is split into two components, one that looks at the possible biases in the method and one that describes the statistical limitations of the method. The first of these two is correlated among the years while the second is uncorrelated.
- The **lepton identification** criteria described in Section 6.5 might have slightly different effects in data and simulation. In order to account for these differences, the efficiency for light leptons to pass the requirements is measured using a “tag-and-probe” method in a Z-enriched region in both data and MC. Scale factors are then applied to simulation to correct for those differences. These corrections come with their own uncertainties on the efficiency measurement and the extrapolation to the signal regions. Because signal leptons are often expected to be quite high in p_T , these identification efficiencies were measured separately in events with reconstructed Z p_T above and below 80 GeV. An uncertainty on these efficiencies based on the difference between these two regions

is then applied as well. Similarly, such corrections were produced for τ_h candidates with p_T below 100 GeV in events with a $\mu\tau_h$ final state. When the $\tau_h p_T$ is below 60 GeV, this region is enriched in DY events while above it the region is designed to be $t\bar{t}$ enriched. Above 100 GeV, single τ_h events in a region targetting highly virtual W bosons are used to derive corrections.

- It was mentioned before that corrections to the **jet energy scale and resolution** are applied. These of course also come with their own uncertainty. In practice, these are obtained for the JEC by varying the jet energy up and down for each jet and propagating the changes in the event through the analysis setup. This will result in different selections and different values for kinematic variables such as p_T^{miss} which will end up in a shape uncertainty on the JEC. The same thing is done for the JER by varying the energy of the smeared jets.
- The **b-tagging** algorithm and its working points come with centrally provided scale factors that correct simulation to match the efficiency in data. Along with these scale factors, up and down variations are provided from which the uncertainty band is obtained by replacing the nominal weight by one where the b-tagging scale factor is replaced by one of the variations. This is done separately with variations of light jet fakes and heavy jets.
- The uncertainty on the data-driven **nonprompt background** estimation technique is derived from the level of closure in the validation results. A separate flat uncertainty with a value of 30% is included for the light lepton nonprompt background, the nonprompt τ_h contribution from DY processes and the nonprompt τ_h contribution from $t\bar{t}$ production. This value is obtained from a study of the level of agreement in the closure tests. They are treated as correlated among all years but uncorrelated relative to each other.
- **Prompt backgrounds** estimated from simulation get assigned a flat uncertainty. For the ZZ and conversion processes the size of this systematic is 10%, estimated from the propagation of the other systematics involved in the control region. They are further constrained by including their control regions in the fit. The same flat uncertainty of 10% is applied to WZ events but on top of that uncertainties on the corrections to the modelling of the M_T^W tail are applied. These tail uncertainties are shape uncertainties estimated from the difference between data and simulation in the $W\gamma$ control region as a function of M_T^W and p_T^{miss} . Their size ranges from 10% at low M_T^W and low p_T^{miss} to 20% at high p_T^{miss} . All other prompt backgrounds receive a flat uncertainty that reflects the precision of theoretical predictions or experimental measurements.
- Our knowledge of the **proton PDFs**, which describe the probability that a certain parton carries a certain fraction of the total proton momentum, is limited and this can be reflected in the modelling in simulation. The uncertainty on these PDFs is estimated by replacing the nominal models by a copy where parameters are varied [173]. This leads to a total of more than 100 different PDFs, of which the root mean square is taken to be applied as the uncertainty shape for this systematic. Similarly, there are

two other aspects of QCD modelling that come with their own uncertainty, namely the **factorisation and renormalisation scales**. Their uncertainty is obtained by varying them simultaneously up and/or down by a factor two and using the envelope as a shape.

An overview of all systematics, their size, the samples they are applied on and their correlation scheme are given in Table 6.8.

6.9 Results

Now that we have all the ingredients in hand, we can unblind⁴ and look at the observed and predicted yields in all discussed search regions. All predicted background yields are obtained as described in Section 6.7 and the total uncertainty includes all systematic uncertainties as described in Section 6.8. The final results can be found in Figures 6.12 to 6.14.

In all search region bins, it can be observed that the data agrees with the background prediction in the SM hypothesis. This can also be observed in Figure 6.15, which shows a comparison between the observed test statistic and the expected test statistic distribution for a background-only fit to the data. When there is agreement, the observed test statistic is expected to be within the bulk of the expected distributions. It should be noted that also in the search channels that were not described in this chapter, no excess of events was observed.

From the results in Figures 6.12 to 6.14, it can be seen that the nonprompt background is the dominant one in most bins and search channels that include a τ_h . It shows the importance of the nonprompt τ_h estimation work performed in the analysis. In all bins where the nonprompt background is the main background, the data is found to agree with the prediction well within uncertainty, both for bins where signal is expected and those where no signal is expected. While this unfortunately means no signal is observed, it also shows the robustness and good performance of the nonprompt estimation methods when used in the signal region.

⁴During the development of an analysis, we do not look at the observed data in the signal regions until all parts of the analysis are in place. This procedure is called blinding and aims to remove bias in the development process.

Table 6.8: An overview of the systematics in the analysis, their size, what processes they apply to and how they are correlated across years [7].

| Source | Typical uncertainty (%) | Processes applied on | Correlation across data-taking years |
|--|-------------------------|--------------------------|--------------------------------------|
| e/ μ efficiency | 1–2 per lepton | Background and signal MC | Correlated |
| τ_h efficiency | 1–3 per lepton | Background and signal MC | Correlated |
| Pileup | 1–2 | Background and signal MC | Correlated |
| Integrated luminosity | 1.8 | Background and signal MC | Partially correlated |
| Trigger efficiency | 1.4–5 | Background and signal MC | Partially correlated |
| Jet energy corrections | 1 | Background MC | Partially correlated |
| Jet energy corrections (fast simulation) | 1 | Signal MC | Correlated |
| b-tagging efficiency | 1–3 | Background MC | Correlated |
| b-tagging efficiency (fast simulation) | 1–3 | Signal MC | Correlated |
| PDF | 1–10 | Signal MC | Correlated |
| Renormalisation and factorisation scales | 1–10 | Signal MC | Correlated |
| Signal initial state radiation (ISR) | 1–5 | Signal MC | Correlated |
| Signal p_T^{miss} | 1–2 | Signal MC | Correlated |
| WZ shape | 5–30 | WZ MC | Uncorrelated |
| WZ normalisation | 10 | WZ MC | Correlated |
| ZZ normalisation | 10 | ZZ MC | Correlated |
| Conversion normalisation | 10 | Conversion MC | Correlated |
| Nonprompt normalisation | 30 | Sideband data | Correlated |
| t \bar{t} X normalisation | 15 | t \bar{t} X MC | Correlated |
| Multiboson normalisation | 50 | Multiboson MC | Correlated |

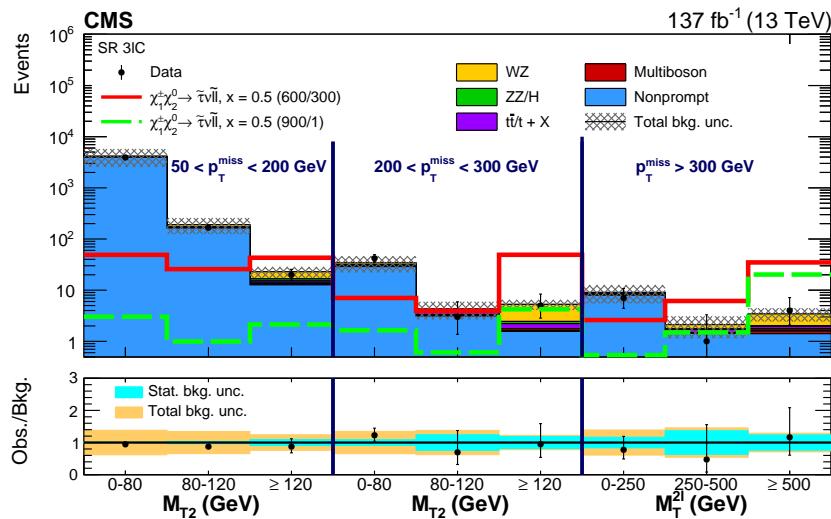


Figure 6.12: Observed and predicted event yields in the search region bins for events with two OSSF light leptons and a single τ_h (3 ℓ C). The solid red and dashed green lines show two models of $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production with τ -enriched slepton-mediated decays in scenarios with $\delta m = 300$ GeV and $\delta m = 900$ GeV respectively, both at $x = 0.5$. The masses of the superpartners for these two hypotheses are given in the legend as $(m_{\tilde{\chi}_2^0}/m_{\tilde{\chi}_1^0})$ in units of GeV. Figure from Ref. [7]

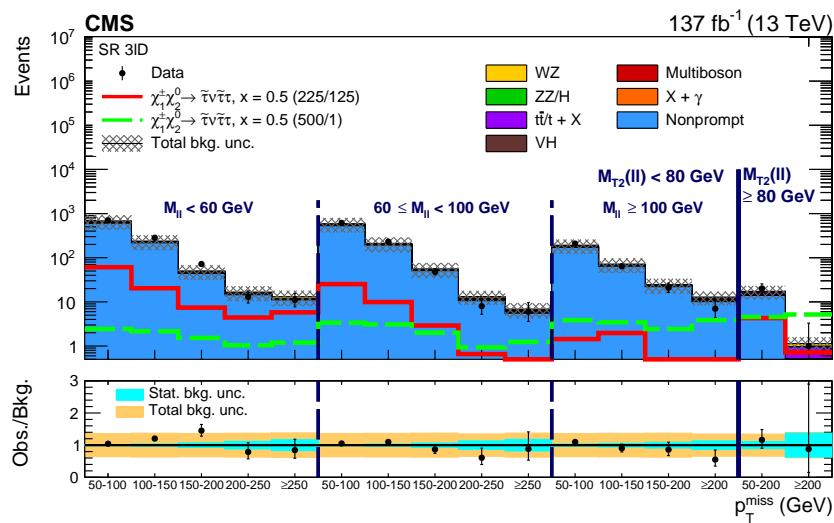


Figure 6.13: Observed and predicted event yields in the search region bins for events with an e^\pm pair and a single τ_h (3ℓD). The solid red and dashed green lines show two models of $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production with τ -dominated slepton-mediated decays in a compressed scenario with $\delta m = 100$ GeV and an uncompressed scenario with $\delta m = 500$ GeV respectively, both at $x = 0.5$. The masses of the superpartners for these two hypotheses are given in the legend as $(m_{\tilde{\chi}_2^0}/m_{\tilde{\chi}_1^0})$ in units of GeV. Figure from Ref. [7]

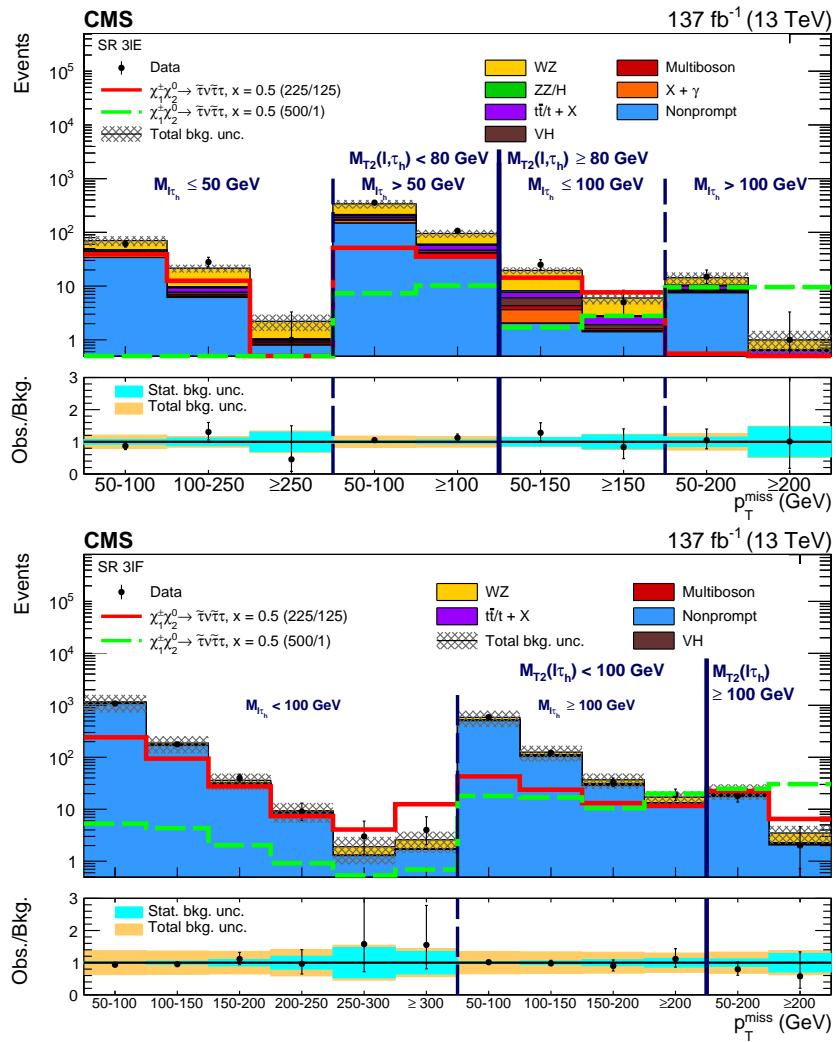


Figure 6.14: Observed and predicted event yields in the search region bins for events with a same-sign light lepton pair and a single τ_h (3 ℓ E) and one light lepton with two τ_h candidates (3 ℓ F). The solid red and dashed green lines show two models of $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production with τ -dominated slepton-mediated decays in scenarios with $\delta m = 100$ GeV and $\delta m = 500$ GeV respectively, both at $x = 0.5$. The masses of the superpartners for these two hypotheses are given in the legend as $(m_{\tilde{\chi}_2^0}/m_{\tilde{\chi}_1^\pm})$ in units of GeV. Figure from Ref. [7]

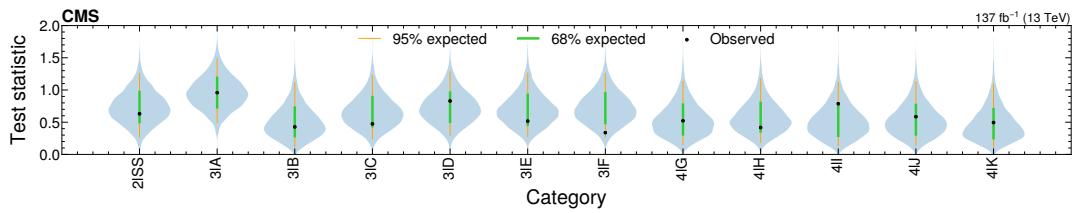


Figure 6.15: Comparison between the observed test statistic value (black dots) and the expected test statistic distribution for a background-only fit for the search regions in the different event categories. The shaded area shows the probability density of the expected test statistic distribution with a green and an orange line that show the 68% and 95% expected range respectively. Figure from Ref. [7]

6.10 Interpretation

The unblinded results show good agreement between data and the SM-only hypothesis. The next logical step in line is to establish exclusion limits on the studied benchmark models. A simultaneous fit is performed to the observed and predicted yields in all relevant search region and control region bins to obtain a 95% CL upper limit on the production cross section of the signal process in question. In this fit, the three years of data-taking are treated as separate inputs. If the theoretical production cross section is higher than this upper limit, it is referred to as excluded at 95% CL. Here we make use of the CL_s criterion [174, 175] with the asymptotic approximation of the distribution of the profile likelihood test statistic [139, 140]. The systematic uncertainties discussed in Section 6.8 are included as nuisance parameters to the fit. The aforementioned flat uncertainties are included in log-normal probability density functions. Shape uncertainties are included via the template morphing technique [176] and assume a Gaussian distribution. The statistical uncertainties of the included samples are taken into account using the Barlow-Beeston approach [177].

For each signal type, the sensitivity in the fit will be driven by different signal regions. For $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production with slepton-mediated decay, both the τ -enriched and τ -dominated scenarios get their main bulk of sensitivity from the trilepton regions. In the case of τ -enriched scenario, the main sensitivity comes from the category with an OSSF light lepton pair and a τ_h (3lC) with smaller contributions from the other channels. It should be mentioned that also the three light lepton final state, which was not discussed in this chapter, contributes non-negligibly due to signal events with leptonic decay of the τ lepton. In tau-dominated scenarios, where we expect higher τ activity, the final state with two τ_h candidates provides the dominant contribution to the sensitivity. In this scenario the contribution from channels not discussed in this chapter is negligible.

The obtained limits for $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production models with slepton-mediated decay in the tau-enriched assumption are shown in Figure 6.16. The coloured blue and red lines show the expected exclusion limits as a function of $m_{\tilde{\chi}_2^0}$ and $m_{\tilde{\chi}_1^0}$ while the black line shows the observed exclusion limit. The green line shows the exclusion limits from the predecessor analysis [178]. An improvement in the limits compared to the previous analysis as a function of $m_{\tilde{\chi}_2^0}$ can be seen of up to 150 GeV and 100 GeV as a function of $m_{\tilde{\chi}_1^0}$. Exclusion limits for τ -dominated $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production on the other hand can be found in Figure 6.17. Here, improvements in the limits of up to 350 GeV in chargino mass and up to 200 GeV in LSP mass are observed. It is clear that in the compressed scenarios, the performance considerably degrades compared to the tau-enriched model. This is due to the high p_T thresholds needed for τ_h reconstruction and identification, which greatly affects the acceptance.

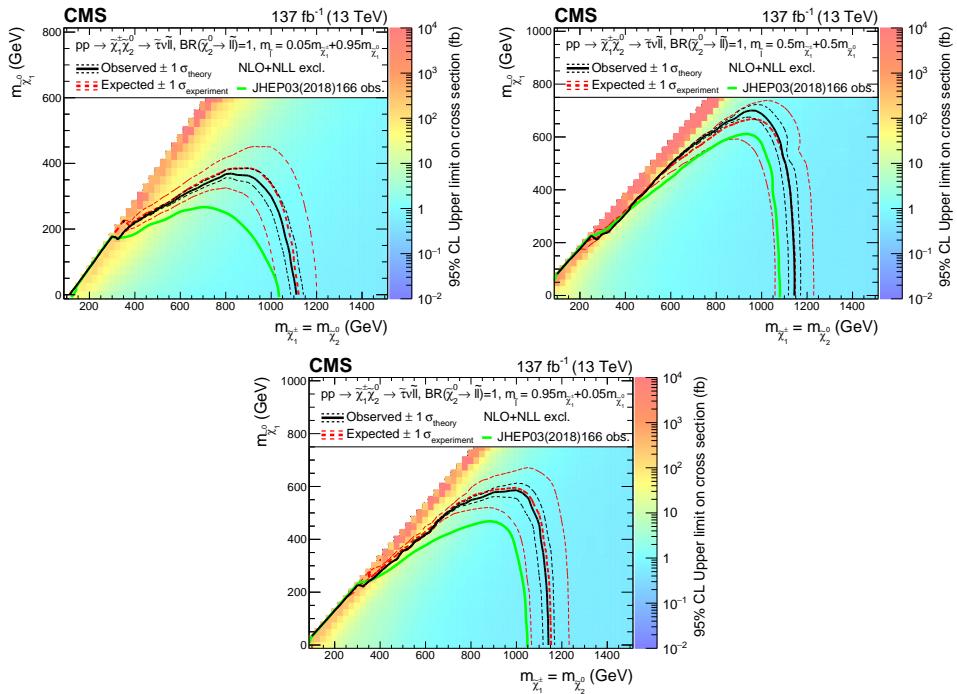


Figure 6.16: Interpretation of the results for $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production in the tau-enriched scenario for three different values of the mass parameter x . The shaded colour represents the 95% CL upper limit on the $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production cross section. The contours show the mass regions that are excluded at 95% CL when assuming the cross section computer at NLO plus NLL. The black contour shows the observed exclusion limits along with the $\pm 1\sigma_{\text{theory}}$, meaning the standard deviation of the theoretical cross sections. The red line shows the expected limit and its $\pm 1\sigma_{\text{experiment}}$ band. The green contour shows the previous CMS limits using 2016 data [178]. Figure from Ref. [7]

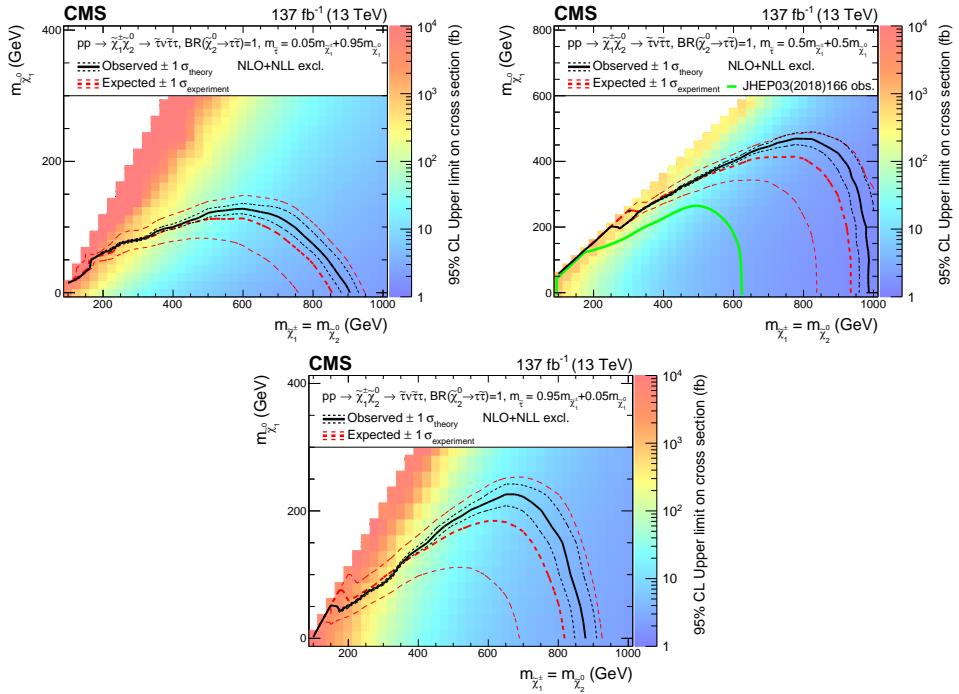


Figure 6.17: Interpretation of the results for $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production in the tau-dominated scenario for three different values of the mass parameter x . The shaded colour represents the 95% CL upper limit on the $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production cross section. The contours show the mass regions that are excluded at 95% CL when assuming the cross section computer at NLO plus NLL. The black contour shows the observed exclusion limits along with the $\pm 1\sigma_{\text{theory}}$, meaning the standard deviation of the theoretical cross sections. The red line shows the expected limit and its $\pm 1\sigma_{\text{experiment}}$ band. The green contour shows the previous CMS limits using 2016 data [178]. Figure from Ref. [7]

6.11 Summary

This chapter described part of an analysis searching for the presence of SUSY in multilepton final states. This was done by analysing data from proton-proton collisions at a centre-of-mass energy of 13 TeV recorded in 2016-2018, corresponding to a total integrated luminosity of 137 fb^{-1} . Search categories were formed by dividing the events up along the number of leptons, their signs and flavours and were then further binned into search regions using discriminating variables that produced phase space regions of high sensitivity to the signal processes. Events were further selected to contain exactly three leptons of which at least one was a τ_h . Techniques for nonprompt background estimation, which forms the dominant background in these channels, are developed and optimised.

An agreement of the data with a SM-only hypothesis was observed. The data was then interpreted to provide upper limits on the production cross section of two variations of signal models describing $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production in a slepton-mediated channel. These models assume that either the $\tilde{\chi}_1^\pm$ decays exclusively to τ leptons (tau-enriched) or that both the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ decay exclusively to τ leptons (tau-dominated). For the $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production with tau-enriched slepton-mediated decay, the observed data translated to 95% CL exclusion limits on the chargino mass up to 1150 GeV. For the second scenario with tau-dominated decays, such limits exclude chargino masses up to 970 GeV.

To finish this chapter, let us briefly discuss the other results of the analysis that fell outside of the scope of this thesis. It should be noted that this analysis has a very high sensitivity to $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ production in flavour-democratic slepton-mediated scenarios as well as scenarios where there is direct decay to a W boson and a Z boson. This is obtained through the development of parametric neural networks where the mass difference between $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ serves as an additional parameter to the training. The inclusion of this mass difference as a parameter allows for a single training to accurately make predictions for a large range of kinematics as well as make accurate interpolations to mass points for which the network was not trained. This strategy allowed us to push forward the $m_{\tilde{\chi}_2^0}$ limit by 400 GeV and the $m_{\tilde{\chi}_1^0}$ limit by 500 GeV in the flavour-democratic slepton-mediated scenario as compared to the predecessor analysis.

These carefully designed upgrades in the analysis strategy and the novel techniques made the analysis so effective. Upon publication, it provided the most stringent limits for chargino-neutralino production in the hardest phase-space to probe, namely for δm close to the Z boson mass. The analysis shines in all slepton-mediated $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ models, providing the most stringent limits in all considered decay modes. The biggest improvements were found in the flavour-democratic scenario where the limit on the chargino mass of 1450 GeV. At the time of publication, this was the highest exclusion value obtained for the production of electroweak superpartners.

Personal contributions

My contributions to this analysis are found in the τ_h sector. I collected and compared all available τ_h identification algorithms at the time and optimised the final τ_h selection of the analysis. I also studied the background due to nonprompt τ_h by measuring the fake rates and performing the closure tests in both simulation and data. Outside of technical analysis work, I have written the documentation on τ_h identification and nonprompt τ_h in the analysis note and answered all questions and comments on the subject during the internal CMS review. Additionally, I have presented an update presentation on this analysis at CMS week and gave a conference talk on the topic of electroweak searches for SUSY at the EPS conference in 2021, which included the analysis discussed in this chapter.

7

A search for heavy neutral leptons in a multilepton final state

7.1 Personal foreword

After exploring the world of SUSY searches guided by the expertise of Willem Verbeke, the time came to start my own analysis. In the summer of 2019, I performed the foundational studies for our search for HNLs in multilepton final states so we could continue to build on them at the start of 2020. This analysis initially envisioned probing the coupling between ν_τ and an HNL, but this scope was broadened to also encompass the light flavour couplings as searched in the predecessor analysis that used only the 2016 dataset. As I was the main analyser of this analysis and it constituted the largest fraction of my PhD, we will go into more detail than we did for the last analysis.

7.2 Introduction

Chapters 2 and 3 have aimed to convey the message that heavy neutrinos can be a solution to multiple problems in the SM. The most prominent of these was the inexplicably small masses of the SM neutrinos, but also the matter-antimatter asymmetry observed in our universe and the question of what constitutes dark matter could be addressed by these particles. The search described in this chapter looks for signs that point at the presence of HNLs in a dataset collected by the CMS collaboration during the Run II period of the LHC. Benchmark models with a single Type-I Seesaw HNL – a singlet under the SM gauge group– of Majorana or Dirac type that couples exclusively to one SM neutrino flavour at a time are considered. In these models, the mass of the HNL m_N and the mixing parameter $|V_{\ell N}|^2$ are the most relevant parameters. Since the HNL is sterile, it will not undergo gauge interactions with SM particles and its interplay with known particles will depend on their mixing with SM neutrinos. Such models have been studied in depth and their signatures in proton-proton predictions have been well predicted [74, 94, 179–183].

This search targets promptly decaying HNLs and studies three different coupling scenarios in final states with three leptons of which at most one can be a τ_h . Events with more than one

τ_h are not included due to low signal acceptance and enormous background yields. For this search, HNL masses in the range of 10 GeV to 1.5 TeV are considered. These are split into two mass categories, one with HNL masses below the W mass and one with sterile neutrino masses above the W mass. Both categories are treated orthogonally. The results in this chapter supersede a similar search performed using only the 2016 dataset [99]. The Run II analysis presented here not only updates Ref. [99] with a larger dataset, but also with improved lepton identification, background estimation techniques and an updated search strategy based on machine learning. It also updates the final states with the inclusion of a τ_h , making it the first result that probes exclusive τ neutrino coupling to HNLs at the LHC.

In this chapter, we will first get into the signal models and their properties in Section 7.3. A brief overview of the data and background samples employed will be given in Section 7.4. This will be followed by the different ingredients needed for the analysis, starting with object selection in Section 7.5 and then with trigger strategy in Section 7.6. Next comes the body of the analysis in the form of event selection and search strategy as described in Section 7.7. Section 7.8 will then give an overview of the main backgrounds and their estimation methods. Finally, all sources of uncertainties are described in Section 7.9 before moving on to the results, interpretation and discussion in Sections 7.10 and 7.11.

7.3 Signal model

The process under study in this analysis consists of the production of a W boson which in turn decays into an HNL. This HNL can decay into either a W boson and a charged lepton or a Z boson and a SM neutrino. When the SM boson decays leptonically, this will end up in a final state of three charged leptons and a SM neutrino. As the HNL to Z decay is subject to a large background of SM processes with a Z boson, the focus will lie on the W decay mode. The corresponding Feynman diagrams in the dominant production mode for masses up to 600 GeV can be found in Figure 7.1. In the case of Dirac HNLs, the charge of the lepton produced directly from HNL decay will need to have opposite charge to the one produced in the initial W decay. For a Majorana HNL, this can both be same-sign or opposite-sign.

It should be noted that the production mode as shown in Figure 7.1, which we will call the DY production mode from now on, is not the only one. While it is by far the dominant process at the lower HNL masses considered in this analysis, up to a few 100 GeV, a second type of production called VBF grows in importance for HNL masses above this range, as illustrated in Figure 7.2 [184, 185]. This latter production mode is shown along with the DY production in Figure 7.3. It is characterized by the presence of an additional quark in the final state that will end up forming a jet.

As mentioned in the introduction, we limit ourselves to benchmark models where only one of the coupling matrix elements has a non-zero value, meaning that the HNL will only couple to one flavour at a time. For each vertex in the Feynman diagram directly linked to the HNL, this will imply that the charged lepton can only have flavour equal to the generation for which

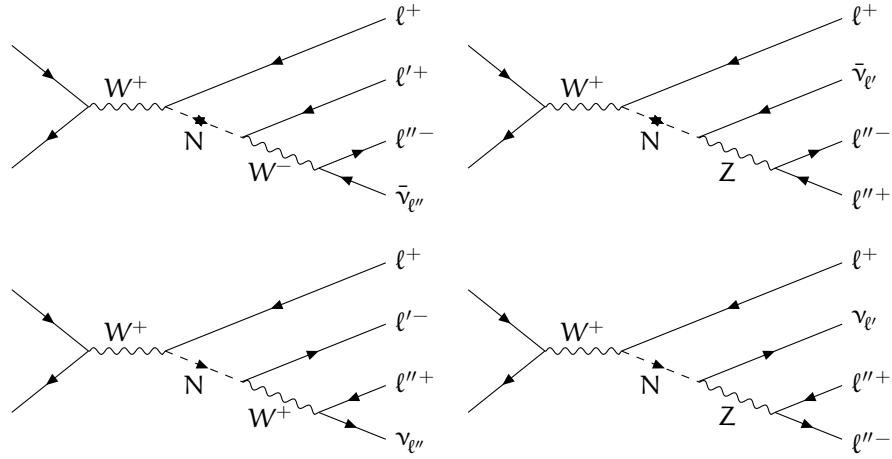


Figure 7.1: Example diagrams for the production of an HNL through the decay of a W boson, in final states with three charged leptons and one SM neutrino. The processes in the upper row are only possible for an HNL of Majorana nature. The processes in the lower row can be realized with an HNL both of Dirac or Majorana nature.

the coupling has a non-zero value. Take for example that $|V_{\mu N}|^2 \neq 0$. The process where the HNL decays to a W boson necessarily must then have two muons, namely from the first two charged leptons. In more complex models where multiple $V_{\ell N}$ entries can be non-zero, this flavour conservation is not present.

While these models where only one flavour couples to the HNL are rather simplistic, they serve as benchmarks for the analysis that are well understood and allow for analysis design that highlights the different coupling aspects. Later re-interpretation techniques can be used to interpret the obtained results in the context of realistic coupling scenarios [186, 187].

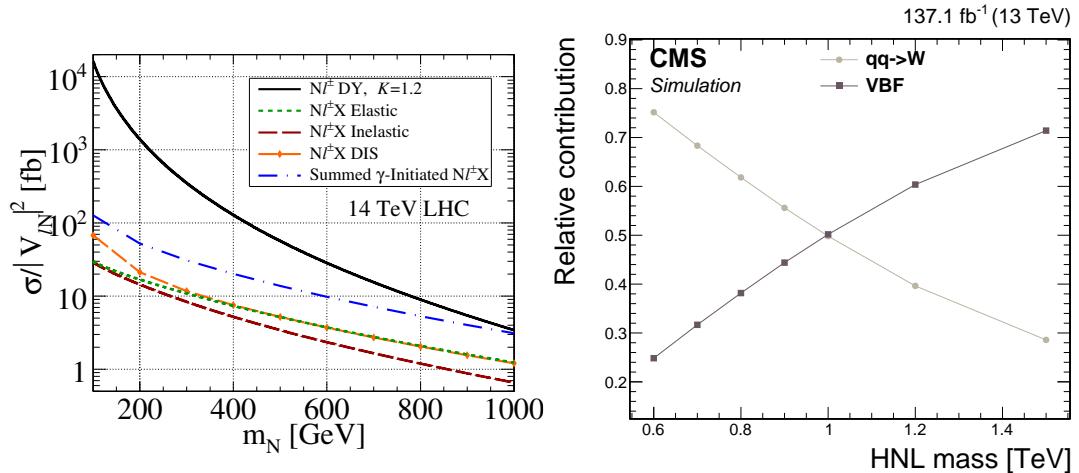


Figure 7.2: Normalized cross sections of the DY (black) and VBF production (blue) channels as a function of mass [184](left) and the relative contribution of the two channels in the final selection of the analysis as measured in the simulated samples with exclusive coupling to electron neutrinos (right).

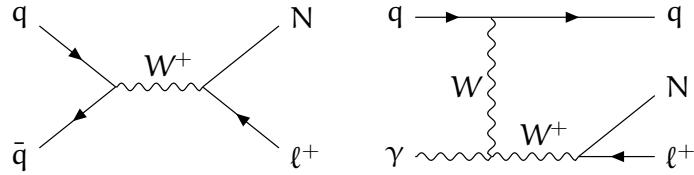


Figure 7.3: Example diagrams for the production of an HNL via DY (left) and VBF (right).

7.3.1 Simulated samples

It is paramount to have signal samples that accurately describe the HNL processes in a search for HNL. The samples for DY and VBF production that were produced for use in this analysis are done so using the `MADGRAPH5_aMC@NLO` generator [157] using the `HeavyN` model [94, 182, 184, 185, 188]. The model itself allows for up to three right-handed Majorana neutrinos that are singlets under the SM gauge groups. In this analysis, we constrain ourselves to samples where there is a single such neutrino that couples to only one lepton flavour at the same time. All three lepton flavours are considered for this exclusive coupling. The time of flight information for the HNLs is not run in these samples, meaning they will all decay promptly in the simulation. In our case, the value of the coupling is set to the value of 0.01:

$$\mathbf{V}_{\ell N} = \begin{pmatrix} 0.01 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{V}_{\ell N} = \begin{pmatrix} 0 & 0 & 0 \\ 0.01 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{V}_{\ell N} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0.01 & 0 & 0 \end{pmatrix}, \quad (7.1)$$

with $\mathbf{V}_{\ell N}$ being the active sterile neutrino mixing matrix:

$$\begin{pmatrix} v_e \\ v_\mu \\ v_\tau \end{pmatrix} = \mathbf{V}_{\ell N} \cdot \begin{pmatrix} N_1 \\ N_2 \\ N_3 \end{pmatrix}. \quad (7.2)$$

While the value of the coupling strength is set to a fixed value, we can use these samples to emulate other coupling values. The reason for this is that the production cross section and the coupling strength squared are proportional to each other. By simply scaling the cross section up or down, we can study different coupling strengths and we do not need to produce a new sample for every coupling we want to probe. The mass of the HNL is a different story, however. As this parameter affects both cross section and event kinematics, we need to produce a new sample for every mass point. For samples where the active coupling to a light lepton is activated, the mass range generated runs from 10 GeV to 1500 GeV in steps of 10 GeV for masses below the W boson mass, 25 GeV for masses between the W boson mass and 200 GeV, 50 GeV for masses between 200 and 500 GeV and steps of 100 GeV above a mass of 500 GeV. This range is slightly shorter for tau coupling, where we only go up to 1000 GeV. At masses above this value, the cross section becomes too small for physics interpretation in this analysis. These discussed ranges were valid for the DY production channels. Because from masses of a few 100 GeV onwards, the VBF channels become more

important, additional samples in this production channel have been generated starting from 600 GeV mass points and above. The generation by **MADGRAPH5_aMC@NLO** was done at LO for models where the HNL has a mass lower than 80 GeV and at NLO for all other mass points [188]. The **NNPDF3.1** sets [189] were the PDFs of choice for all simulated samples. The hadronization, parton showering and underlying event simulation is taken care of by **PYTHIA v8.230** [190], using the CP5 tune [191] for all 4 datasets. Double counting between **PYTHIA** and **MADGRAPH5_aMC@NLO** is removed using the MLM matching scheme [192] for LO generation and the FxFx matching scheme [193] for NLO generation.

In order to increase the statistics in the final selection, filters on the p_T of generated particles were applied in the **PYTHIA** step of the sample generation. These filters aim to bring the generated p_T thresholds close to the p_T thresholds applied at the analysis level. By doing this, events that would be discarded by the analysis thresholds are already removed in the first steps of sample generation and do not have to go through the more time-intensive following steps of generation. This means that these samples have a higher acceptance in the final selection than samples without such a filter with the same number of total events. The first filter targets fully light lepton final states and requires there to be at least three light leptons with a p_T larger than 5 GeV in the generated event. The second filter targets final states containing a τ_h and thus requires the presence of at least one τ_h with visible p_T – meaning not taking the neutrinos from the τ_h decay into account – larger than 18 GeV. This reflects the fact that offline τ_h reconstruction starts from a p_T requirement of 20 GeV and takes into account that reconstruction will not be perfect by using a 2 GeV buffer that ensures we do not lose acceptance. Additionally, this filter requires at least one light lepton with $p_T > 15$ GeV to reflect trigger requirements, as will become more clear in Chapter 7.6. The light lepton filter is applied in all coupling scenarios with exclusive coupling to light flavour neutrinos. As both light lepton final states and final states containing a τ_h are used to study scenarios with exclusive coupling to τ -neutrinos, two sample types were generated. One of these uses the light lepton filter and the other uses the τ_h filter.

Because the search looks for prompt HNL decay, the bulk of our signal samples is prompt, meaning the lifetime of the HNL is neglected. At the lowest HNL masses, there is some displacement to be expected. In the simulated signal process with an HNL mass of 10 GeV and $|V_{eN}|^2 = 10^{-5}$, for example, an average lifetime of the HNL of about 0.4 mm is observed¹, which can translate to displacements of a few cm in the CMS frame. We will have to take this into account when evaluating our final yields. Likewise, all samples are generated assuming the neutrinos are Majorana in nature. In the end, we do not want to constrain ourselves to this single assumption and also provide an interpretation of models with a Dirac HNL. Both of these aspects are briefly discussed below.

¹The average lifetime is expressed here in a length scale as $c\tau$ where c is the speed of light and τ is the average lifetime of the HNL in its own rest frame in seconds.

Taking into account displacement

Due to the relation between HNL mass and lifetime, when one goes below masses of 20 GeV, the assumption that there will be no displacement in the HNL decay breaks down. Because the prompt samples ignore HNL time of flight, there will be an overestimation of events that pass the selection. The reason for this is that the displaced leptons would in reality not pass the object selection where quality cuts on the impact parameters are made. We solve this by producing more samples. This time, however, we take into account proper displacement and do this only for samples with a mass of 20 GeV or lower. At this point, the flexibility of being able to probe different coupling strengths by simply scaling the cross section is lost. For a fixed mass, going to a different coupling strength means going to a different average lifetime which will affect the signal selection efficiencies and thus the signal yield. At first, it looks like a different sample would have to be generated for every coupling strength value. This is not feasible when it comes to limit production as we would have to iteratively produce new samples to inch closer to a signal strength of one (see Section 7.10.1).

Instead, we use a reweighting strategy using the analytical form of the *decay probability distribution*:

$$\frac{dN(t)}{dt} = \frac{1}{\tau_N} e^{-t/\tau_N}, \quad (7.3)$$

which determines the probability distribution for HNLs with *mean lifetime* τ_N to have a *proper lifetime* t as measured in the HNL rest frame. Because the average kinematics of the event will not significantly depend on the mean lifetime but only on the HNL mass and momentum, it is perfectly possible to emulate a specific target mean lifetime τ_{target} by reweighting events from a sample with the same HNL mass but different mean lifetime τ_0 by using the proper ratio of probability distribution values for the proper lifetime of the HNL:

$$T(t) = \frac{dN_{target}(t)/dt}{dN_0(t)/dt} = \frac{\tau_0}{\tau_{target}} \frac{e^{-t/\tau_{target}}}{e^{-t/\tau_0}}. \quad (7.4)$$

This means that in order to probe HNL scenarios with the same mass but different coupling squared values we will have to weight each event in the original sample with the usual cross section reweighting – the ratio of the coupling squared – and the lifetime weight $T(t)$.

Probing Dirac scenarios

The practical difference between Majorana and Dirac scenarios is that the first allows for lepton number violating events where the lepton from original W-production and the one that is directly produced in the HNL decay have the same sign, meaning they are both leptons or anti-leptons. Dirac-type HNL can not allow this violation and will result in events where the first two leptons have opposite sign. Taking a closer look at this difference, it means that a Majorana HNL has twice the amount of decay channels than a Dirac HNL with the same mass and coupling values, which also means it has exactly twice the resonance width and half the lifetime.

With the exception of the lifetime – for which there is already a reweighting procedure in place – this difference is only seen in final state combinatorics. It is possible to obtain information on Dirac-type HNLs from our already existing Majorana samples. For prompt samples, this simply comes down to checking all events in the original, unskimmed sample, identifying the lepton from initial W decay and the light neutrino from HNL decay on generator level, requiring that they have opposite lepton number and adding them to a counter to be used in luminosity scaling. This latter counter of the total number of opposite-sign events should be around a factor of 0.5 of the total number of events. For displaced samples the procedure is identical but this time a factor of 0.5 is needed in the exponents of the lifetime correction $T(t)$ to account for the decreased mean lifetime.

7.3.2 Signal compression

One important realization to make before designing an analysis to probe for processes of the type shown in Figure 7.3, is that the event kinematics will be quite different for HNL masses below and above the W boson mass. When the mass of the HNL is smaller than that of a W boson, it can be seen from the Feynman diagram that the W produced in the HNL decay has to be a virtual boson. Additionally, the kinematics will depend on the size of the difference between the HNL and W masses. For mass points such as a 20 GeV HNL where this difference is relatively large, the energy will go into the charged lepton coming directly from the initial W decay while less energy will be available for distribution among the three HNL decay products. In scenarios where this mass difference is rather small – also referred to as compressed –, such as a 60 GeV HNL, the scenario is inverted and the initial charged lepton will be relatively soft.

When inverting the relative mass scale between HNL and W boson, the initial W boson will have to be a virtual one, which has two consequences. Firstly it is the reason why the production cross section for the signal process drops so steeply with increasing m_N once we go higher than the W mass. The production cross section in the DY production mode takes the following form [94, 194]:

$$\begin{aligned} \sigma(pp \rightarrow l_1 l_2 W) &\approx (2 - \delta_{l1l2}) \sigma(pp \rightarrow l_1 N) Br(N \rightarrow l_2 W) \\ &= (2 - \delta_{l1l2}) \frac{|V_{l1N} V_{l2N}|^2}{\sum_{l=e}^{\tau} |V_{\ell N}|^2} \sigma_0(N), \end{aligned} \quad (7.5)$$

where $\sigma_0(N)$ is called the *bare cross section* and is mainly dependent on the W production cross section. This cross section is shown as a function of HNL mass in Figure 7.4. Br denotes the *branching ratio*, which describes what fraction of the decays of a certain kind of particle will occur through that specific decay mode.

Secondly, it implies that the virtual mass of the initial W boson will mostly reside as close to m_W as possible since the cross section will be the highest there [195]. This means that in scenarios where the HNL has a mass that is much higher than the m_W mass, although the

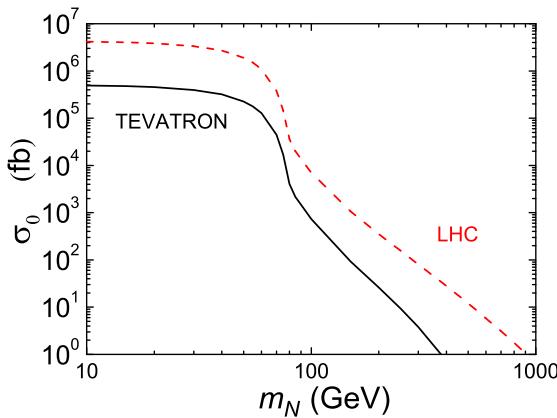


Figure 7.4: The bare cross section for HNL production as a function of HNL mass in different collider assumptions [194].

lepton momenta spectra are typically harder than their low-mass counterparts, the charged lepton from initial W decay tends to be relatively soft compared to the HNL decay products.

In general, the signal process will thus have at least one relatively soft lepton. In compressed scenarios, even the leading lepton might still have a p_T lower than 40 GeV. We need to keep this into account when designing the object selection, trigger selection and search strategy as we need to go as low as possible in p_T to retain signal acceptance.

7.4 Data and background simulation

All data and simulation used in the analysis are part of the *ultra-legacy (UL)* reprocessing campaign of the Run II data-taking period. This reprocessing deals with a detector issue in 2016 due to highly ionizing particles [196]. When there are inelastic interactions between hadrons and the silicon detector, fragments can be produced that saturate the front-end chips temporarily. This causes dead-times that last around 700 ns. The issue was patched in the middle of the 2016 data-taking period by tweaking the settings of the front-end electronics and was addressed in the UL reprocessing by using adapted reconstruction techniques for the affected part of the dataset. At the same time, the UL reprocessing came with updated calibrations and scale factors for the different subdetectors and the reconstruction and identification algorithms. The most prominent of these new calibrations can be found in the ECAL. There, a relative improvement in the electron energy resolution of up to 37% can be found in the regions of highest pseudorapidity [168, 197, 198].

7.4.1 Collision data

We make use of Run II data but, unlike the three years of data-taking in Chapter 6, it is split up into four sets of proton-proton collisions. Two of them are the familiar 2017 and 2018 datasets,

each bringing 41.48 fb^{-1} and 59.83 fb^{-1} of integrated luminosity respectively.² The other two sets correspond to the first 19.5 fb^{-1} of 2016 data where the detector issue was present and the last 16.8 fb^{-1} where the issue was resolved. They will be referred to as pre-VFP and post-VFP 2016 data respectively³. Default reconstruction methods were used for the processing of the post-VFP dataset while adaptations were made in the pre-VFP dataset in order to mitigate the issue. The four sets are processed independently in the analysis and will each have their own set of simulations, calibrations, etc.

7.4.2 Standard model process simulation

Simulated samples have been produced centrally for each of the SM background processes in all four datasets. This was done so in the Summer20UL16, Summer20UL17 and Summer20UL18 campaigns. The `MADGRAPH5_aMC@NLO` [156, 157], `POWHEG v2` [158–161] and `MCFM` [162] programs are used for generation of the samples. More specifically, `POWHEG` was the generator of choice for all $t\bar{t}$, t-channel single top, tW and qq production modes of diboson processes and H boson production. All aforementioned samples were simulated at NLO in the strong coupling constant. Diboson processes where the initial state consists of two gluons were generated using the `MCFM` generator at LO. All other background samples were produced using `MADGRAPH5_aMC@NLO` either at LO or NLO in QCD.

Just like the signal simulation, the chosen PDFs are the `NNPDF3.1` sets [189] for all simulated samples. Also here, the hadronization, parton showering and underlying event simulation are taken care of by `PYTHIA v8.230` [190] using the `CP5` tune [191] for all 4 datasets and double counting between `PYTHIA` and `MADGRAPH5_aMC@NLO` is removed using the `MLM` matching scheme [192] for LO generation and the `FxFx` matching scheme [193] for NLO generation.

7.5 Object Selection

In the same way as in the SUSY search, all particle objects are reconstructed using the PF algorithm and jets are clustered from reconstructed particles using the anti- k_T algorithm with a distance parameter of 0.4 (AK4) [124]. While the baseline selection is very similar, the identification techniques have been updated in this analysis.

7.5.1 Light Leptons

Also here, leptons are the stars of the show. In contrast to the last iteration of the analysis, the decision was made to go for a machine learning-based approach for light lepton identification. The lepton MVAs used are not the same as described in Chapter 6 but are more up-to-date

²You may have noticed that the integrated luminosity for these two years is slightly different than in Chapter 6. This is the case because the UL processing comes with an updated catalogue of certified events that can be used for analysis purposes.

³The name VFP comes from the Pre-amplifier Feedback Voltage Bias parameter which was set to zero in the post-VFP data-taking to fix the issue.

MVAs trained in-house at Ghent University in the context of searches for processes with top quarks [199].⁴ Again a system of three working points to cover baseline selection, sideband selection and final selection is employed.

Lepton MVA

The lepton MVAs employed here are a result of the hard work of Kirill Skovpen, who developed highly efficient BDTs designed for use in light lepton identification. Although it was initially designed for searches involving top quarks, no performance is lost when using them outside of these searches. In order to have a unified strategy across our working group with well-validated and cross-checked discriminators, it was ported for use here. Taking a look behind the scenes, the leptonMVAs consist of gradient BDT trained using the TMVA package [200]. They are trained to discriminate prompt leptons, defined as non-displaced signal-type leptons from W , Z and τ decay, from nonprompt leptons, defined as leptons from background sources such as heavy-flavour jets or jet fakes.

Different MVAs were trained for discrimination against muons and electrons. This is done using input variables that describe the kinematics of the lepton candidate, the isolation in the vicinity of said candidate, impact parameters of the lepton candidate, properties of the closest jet, electron MVA ID (see Section 6.5.2) and compatibility of track segments of the muon in different subdetectors. The first set describes the baseline properties of the lepton such as p_T and η . The second set, namely the isolation parameters, are expected to have small values for prompt leptons while nonprompt leptons tend to have larger hadronic activity in their neighbourhood. The impact parameters help to discriminate non-displaced leptons against displaced leptons while the nearest jet properties give more insight into nonprompts from jets and heavy-flavour quarks. In the latter case, the nearest jet is defined as the jet with which the lepton shares a PF candidate.

We will not go into the details of the architecture here. Instead, we will simply quote the performance of the MVAs in Figure 7.5. The working point that will be used for muons is labelled as Medium and for electrons as Tight on the “TOP-UL” ROC curve in this figure. It can be seen that we reach over 95% per lepton while retaining high background discrimination.

Muons

Muon selection cuts on p_T , $|\eta|$, impact parameters, isolation and muon POG ID in the same way as described in Section 6.5.1. The main difference comes from the use of the different lepton MVA. Additionally, the FO object makes use of selections on the p_T^{ratio} and the b-tagging score of the closest jet ($\text{DEEPFLAVOR}^{\text{closest}}$) for both of which the selection of the closest jet is identical to the way it is done in the training of the MVA. A full overview of all selections for all three working points is given in Table 7.1. All muons have Rochester corrections [167] and identification scale factors applied.

⁴The lepton MVA referred to from Chapter 6 refers to the final lepton MVA used in the tight selection, not the electron MVA.

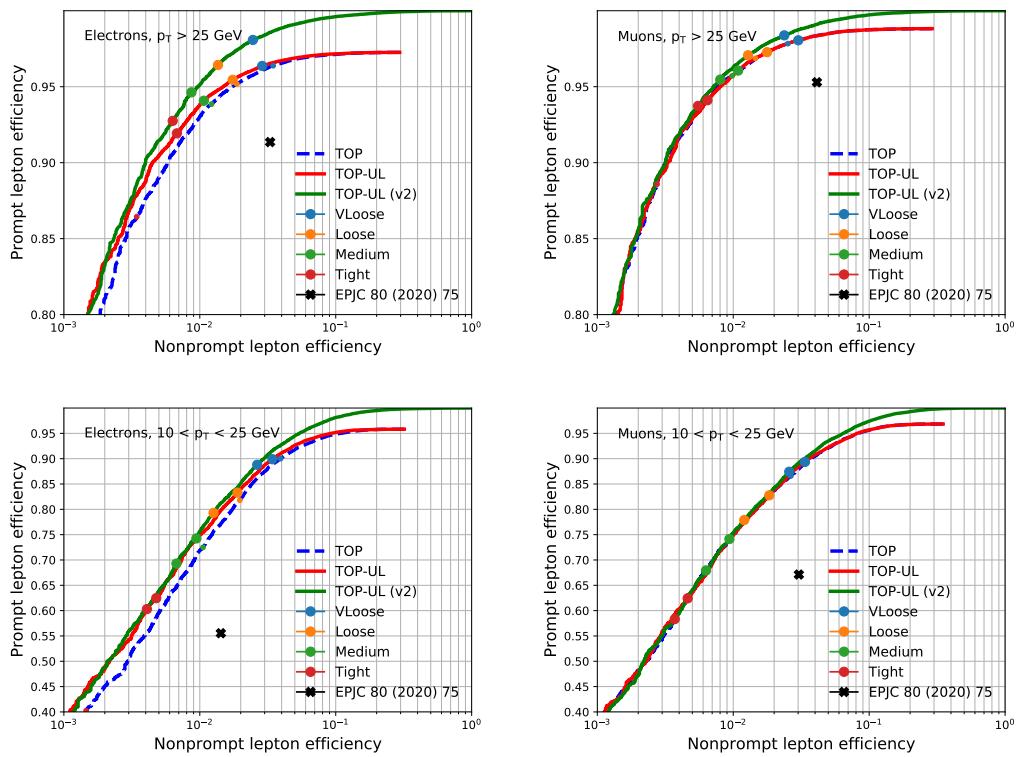


Figure 7.5: Performance of Top Lepton MVA ID in $t\bar{t}$ events for electrons (left) and muons (right), with $p_T > 25$ GeV (top) and $10 < p_T < 25$ GeV (bottom). Plots by Kirill Skovpen.

Electrons

In the SUSY analysis, we came across a lot of variables describing the shape of the cluster and the matching between the ECAL cluster and the track of the electrons. This information is now included in the MVA and was not filtered on in the preselection of its training. For that reason, we also do not include most of these in the electron preselection and let the MVA do its work. A list of all selection cuts that are applied can be found in Table 7.2. On top of these selections, all electrons that have a loose muon within a distance of $\Delta R < 0.05$ are removed from the analysis. Electron scaling and smearing corrections are propagated to all simulated electrons used in the analysis in order to match the energy of the reconstructed electrons in the analysis closer to their true energy, and reconstruction scale factors are applied to all simulated electrons passing the loose working point to correct the reconstruction efficiency in simulation. Additionally, a set of identification scale factors for lepton MVA efficiency were applied.

7.5.2 Tau Leptons

Tau leptons form a new addition to the prompt HNL analysis and were not present in the previous iteration. Their reconstruction is performed using the HPS algorithm we are familiar with at this stage and the baseline identification consists of the regular 20 GeV p_T cut and the requirement that the τ_h is found in the fiducial region of $|\eta| < 2.3$.

At this point in time, the DeepTau algorithm was no longer an experimental project but fully

Table 7.1: Selection requirements that define the muon ID WPs. Cuts that are labelled with * are applied in the case that the muon does not meet the muon MVA requirement.

| Cut | Loose | FO | Tight |
|-----------------------------|-------|----------|-------|
| p_T (GeV) | >5 | >10 | >10 |
| $ \eta < 2.4$ | ✓ | ✓ | ✓ |
| $ d_{xy} < 0.05$ cm | ✓ | ✓ | ✓ |
| $ d_z < 0.1$ cm | ✓ | ✓ | ✓ |
| $I_{\text{mini}} < 0.4$ | ✓ | ✓ | ✓ |
| $\sigma_{\text{IP,3D}} < 8$ | ✓ | ✓ | ✓ |
| PF-medium | ✓ | ✓ | ✓ |
| p_T^{ratio} | — | >0.45 * | — |
| Deepflavour(closest jet) | — | <0.025 * | — |
| lepton MVA > 0.64 | — | — | ✓ |

Table 7.2: Selection requirements that define the electron ID WPs. Cuts that are labelled with * are applied in the case that the electron does not meet the electron MVA requirement. Cuts with multiple defined values refer to the respective cuts for 2016 and 2017/2018. Additionally, all electrons are required to be $\Delta R > 0.05$ away from a loose muon.

| Cut | Loose | FO | Tight |
|--------------------------------|-------|------------|-------|
| p_T (GeV) | >10 | >10 | >10 |
| $ \eta < 2.5$ | ✓ | ✓ | ✓ |
| $ d_{xy} < 0.05$ cm | ✓ | ✓ | ✓ |
| $ d_z < 0.1$ cm | ✓ | ✓ | ✓ |
| $I_{\text{mini}} < 0.4$ | ✓ | ✓ | ✓ |
| $\sigma_{\text{IP,3D}} < 8$ | ✓ | ✓ | ✓ |
| missing hits < 2 | ✓ | ✓ | ✓ |
| Conversion veto | — | ✓ | ✓ |
| electron MVA | — | WP-Loose * | — |
| p_T^{ratio} | — | >0.5/0.4 * | — |
| Deepflavour(closest jet) > 0.1 | — | ✓* | — |
| lepton MVA > 0.81 | — | — | ✓ |

fleshed out and validated. As such, we have used DeepTauv2p1 as the identification algorithm of choice in this analysis. To accompany it, a selection is performed on the decay mode of the reconstructed τ_h candidate, allowing only for decay modes with one or three charged pions and a number of π^0 . The algorithm has taken strides in its identification efficiency and low misidentification rates compared to its predecessor MVA as can be seen in Figure 7.6. In order to keep the balance between low misidentification and high efficiency to get optimal sensitivity in the final state, the medium working point of the DeepTauVSjet output is used to filter τ_h against jets faking their signature. To remove fakes from electrons and muons, the tight working points of the relevant output scores are applied. A final cleaning removes any τ_h from the object list if it is found to be within $\Delta R < 0.5$ of a tight electron or muon. This is done to further reduce fakes and avoid potential double counting.

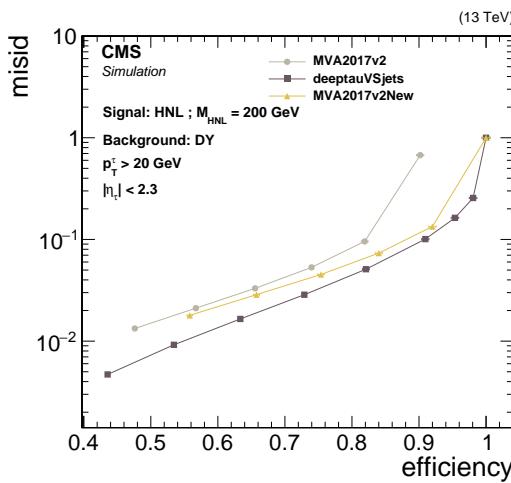


Figure 7.6: Efficiency versus jet misidentification rate of the DeepTau algorithm and two versions of the MVA discriminator. The efficiency is measured in an HNL sample with exclusive τ neutrino coupling while the misidentification rate is measured in a DY sample.

In contrast to the light leptons, it is not easily possible to create a FO selection that works for our baseline selection when using the DeepTau working points. Therefore we will find ways to circumvent the need for this type of working point as we will discuss in Section 7.8.1 and instead stick to a loose and a tight working point. For this loose working point, the selection on the DeepTauVSjet output score is loosened to the loosest working point available (vvvloose). An overview of the selections that define both of them can be found in Table 7.3.

In the same way as in Section 6.5.3, corrections for tau energy scale and identification are applied using the same recipe but this time using the relevant corrections for the DeepTau ID.

7.5.3 Jets

All AK4 jets that want to qualify as useful to the analysis should have a p_T higher than 20 GeV, be within the calorimeter region of $|\eta| < 2.4$, pass the Tight jet ID (see Section 6.5.4) and contain no loose light lepton within $\Delta R < 0.4$ and no tight τ_h within $\Delta R < 0.5$. Also here JEC and JER corrections are applied.

Table 7.3: Selection requirements that define the τ lepton ID WPs.

| Cut | Loose | Tight |
|----------------------|-------------|-----------|
| p_T (GeV) | >20 | >20 |
| $ \eta < 2.3$ | ✓ | ✓ |
| $ d_z < 0.2$ cm | ✓ | ✓ |
| DeepTau VS jets | vvvloose WP | medium WP |
| DeepTau VS electrons | tight WP | tight WP |
| DeepTau VS muons | tight WP | tight WP |

B-tagged jets are selected from this baseline jet selection with the difference that the cleaning is loosened. Instead of the loose lepton working point, b-jets should not reside within $\Delta R < 0.4$ of FO leptons. To perform the tagging, the Deepflavour algorithm – which we recognize from the SUSY analysis – is used to select two working points. A loose selection aims to select b-jets to be used in the b-veto that we will introduce in later sections. The tight working point has only the selection of b-jets in a $t\bar{t}$ control region in mind. B-tagging scale factors are applied to all b-jets in simulation using the same recipe as in Section 6.5.4.

7.5.4 Missing transverse energy

To get a handle on the missing momentum in the events, a missing momentum vector \vec{p}_T^{miss} is defined as the inverted vector sum of the transverse momentum of all PF candidates (PFMET). The missing transverse momentum is then defined as the size of this vector. The previously mentioned JEC will impact the calculation of this quantity and thus these corrections are propagated to the \vec{p}_T^{miss} . In similar fashion as the analysis in Chapter 6, MET filter are applied.

7.6 Trigger strategy

In the search for electroweak SUSY signatures, it was briefly mentioned that all events need to pass a certain combination of triggers. It is important that we apply the same strategy here. As the trigger system forms the basic “garbage disposal” of uninteresting events, we need to make sure the data and simulation start from the same baseline. This is done by requiring that the events we select pass at least one trigger of interest to the final state of this analysis.

7.6.1 Strategy

There are three types of lepton triggers: single lepton triggers, double lepton triggers and trilepton triggers. In order to keep the rate under control, single lepton triggers have higher online p_T thresholds than double and trilepton triggers. Keeping in mind that our signal signature is characterized by soft leptons, the trigger cocktail is designed to contain all

available light lepton triggers with the lowest p_T thresholds that are not prescaled⁵. As such, we end up with a mixture of single light lepton, double light lepton and triple light lepton triggers. This allows to make the p_T selection for the leptons as low as possible. While there are triggers available that select on the presence of a τ_h , as described in Section 5, it was chosen to not use them here. The reason for this is that the online thresholds are quite high and the turn-on curve of the triggers is slow. This makes them very useful for analyses with a hard τ_h spectrum but not as much so here. To use these triggers would mean applying high thresholds on the offline objects, which in turn means that they do not add much efficiency to the total cocktail while at the same time making the trigger analysis more complex. A full overview of all triggers used in this analysis is given for 2016, 2017 and 2018 data separately in Tables 7.4 to 7.6.

Table 7.4: Summary of 2016 HLT triggers grouped by channel.

| Channel | HLT paths |
|-------------|---|
| 3e | HLT_Ele16_Ele12_Ele8_CaloIdL_TrackIdL |
| 2e, 1 μ | HLT_Mu8_DiEle12_CaloIdL_TrackIdL |
| 1e, 2 μ | HLT_DiMu9_Ele9_CaloIdL_TrackIdL |
| 3 μ | HLT_TripleMu_12_10_5 |
| 2e | HLT_Ele23_Ele12_CaloIdL_TrackIdL_IsoVL_DZ |
| 1e, 1 μ | HLT_Mu23_TrkIsoVVL_Ele8_CaloIdL_TrackIdL_IsoVL HLT_Mu8_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL_DZ HLT_Mu8_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL HLT_Mu23_TrkIsoVVL_Ele8_CaloIdL_TrackIdL_IsoVL_DZ |
| 2 μ | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ HLT_Mu17_TrkIsoVVL_TkMu8_TrkIsoVVL HLT_Mu17_TrkIsoVVL_TkMu8_TrkIsoVVL_DZ HLT_TkMu17_TrkIsoVVL_TkMu8_TrkIsoVVL HLT_TkMu17_TrkIsoVVL_TkMu8_TrkIsoVVL_DZ |
| 1e | HLT_Ele27_WPTight_Gsf |
| 1 μ | HLT_IsoMu24 HLT_IsoTkMu24 |

7.6.2 Offline thresholds

The next step in the trigger strategy is then to make sure that all triggers operate in the high-efficiency plateau. A baseline selection of three leptons with $p_T > 15, 10, 10$ GeV covers the events with three muons or two muons and an electron. On top of this baseline selection, we apply an offline threshold strategy based on the leading and subleading light leptons in the event. Events should either have a leading light lepton with a p_T high enough to be in the

⁵Prescaling refers to a practice performed in the triggering system when the rate of a path is too high for the budget. When a prescale is set to i.e. 2, it means that only one out of two events that pass the trigger will be stored.

Table 7.5: Summary of 2017 HLT triggers grouped by channel.

| Channel | HLT paths |
|-------------|--|
| 3e | HLT_Ele16_Ele12_Ele8_CaloIdL_TrackIdL |
| 2e, 1 μ | HLT_Mu8_DiEle12_CaloIdL_TrackIdL |
| 1e, 2 μ | HLT_DiMu9_Ele9_CaloIdL_TrackIdL_DZ |
| 3 μ | HLT_TripleMu_12_10_5 |
| | HLT_TripleMu_10_5_5_DZ |
| 2e | HLT_Ele23_Ele12_CaloIdL_TrackIdL_IsoVL |
| 1e, 1 μ | HLT_Mu23_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL_DZ |
| | HLT_Mu23_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL |
| | HLT_Mu8_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL_DZ |
| | HLT_Mu8_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL |
| 2 μ | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL |
| | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ |
| | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass8 |
| | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass3p8 |
| 1e | HLT_Ele32_WPTight_Gsf |
| 1 μ | HLT_IsoMu24 |
| | HLT_IsoMu24_eta2p1 |

Table 7.6: Summary of 2018 HLT triggers grouped by channel.

| Channel | HLT paths |
|-------------|--|
| 3e | HLT_Ele16_Ele12_Ele8_CaloIdL_TrackIdL |
| 2e, 1 μ | HLT_Mu8_DiEle12_CaloIdL_TrackIdL |
| 1e, 2 μ | HLT_DiMu9_Ele9_CaloIdL_TrackIdL_DZ |
| 3 μ | HLT_TripleMu_12_10_5 |
| | HLT_TripleMu_10_5_5_DZ |
| 2e | HLT_Ele23_Ele12_CaloIdL_TrackIdL_IsoVL |
| 1e, 1 μ | HLT_Mu23_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL_DZ |
| | HLT_Mu23_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL |
| | HLT_Mu8_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL_DZ |
| | HLT_Mu12_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL_DZ |
| 2 μ | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL |
| | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ |
| | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass8 |
| | HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass3p8 |
| 1e | HLT_Ele32_WPTight_Gsf |
| 1 μ | HLT_IsoMu24 |

high-efficiency region of any relevant single light lepton trigger or have leading and subleading light leptons that reside in the plateau of the relevant double light lepton trigger. The way this is done in practice is shown in Table 7.7. Here every row represents offline thresholds that correspond to a single trigger. On the left-hand side, the lepton requirements for that specific trigger are displayed. When an event passes at least one of the rows, it is retained in the analysis.

Table 7.7: Additional requirements on the light lepton p_T depending on the flavour combination of the event to be applied on top of a baseline selection of three leptons with $p_T > 15, 10, 10$ GeV. The columns give the numbers of electrons and muons in the event (N_e and N_μ), and the p_T of the leading electron, subleading electron, and leading muon, respectively. All events are required to pass the conditions of at least one of the rows. The values in parentheses give the threshold applied in 2017 and 2018, where it is different from 2016.

| N_e | N_μ | $p_T(e_1)$ [GeV] | $p_T(e_2)$ [GeV] | $p_T(\mu_1)$ [GeV] |
|----------|----------|------------------|------------------|--------------------|
| ≥ 1 | — | >30 (35) | — | — |
| — | ≥ 1 | — | — | >25 |
| ≥ 1 | ≥ 1 | >25 | — | >10 (15) |
| ≥ 1 | ≥ 1 | >10 (15) | — | >25 |
| ≥ 2 | — | >25 | >15 | — |
| — | ≥ 2 | — | — | >20 |
| $= 1$ | $= 2$ | — | — | — |
| — | $= 3$ | — | — | — |

7.6.3 Trigger efficiency

Trigger efficiency in simulation

In order to check the efficacy of the offline threshold strategy, the total trigger efficiency is measured in simulation as a function of subleading and trailing lepton p_T . In order to perform the measurement, a region of events with three leptons that pass the tight working point and that pass the offline thresholds is selected. The efficiency is then defined as the fraction of events that also pass the triggers. The efficiency for different types of final state and kinematics are shown in Figures 7.7 and 7.8. Apart from a few exceptions, the efficiency is 90% or higher everywhere.

Trigger efficiency in data

Now that the efficiency of the triggers has been validated to be in the high-efficiency regime, a check of this efficiency in data needs to be performed as well. The reconstruction techniques in data and simulation will respond slightly differently and by comparing the efficiency in both, we have a handle on how good the agreement is and whether we need additional corrections. In order to be able to compare data and simulation, we need to perform a preselection of

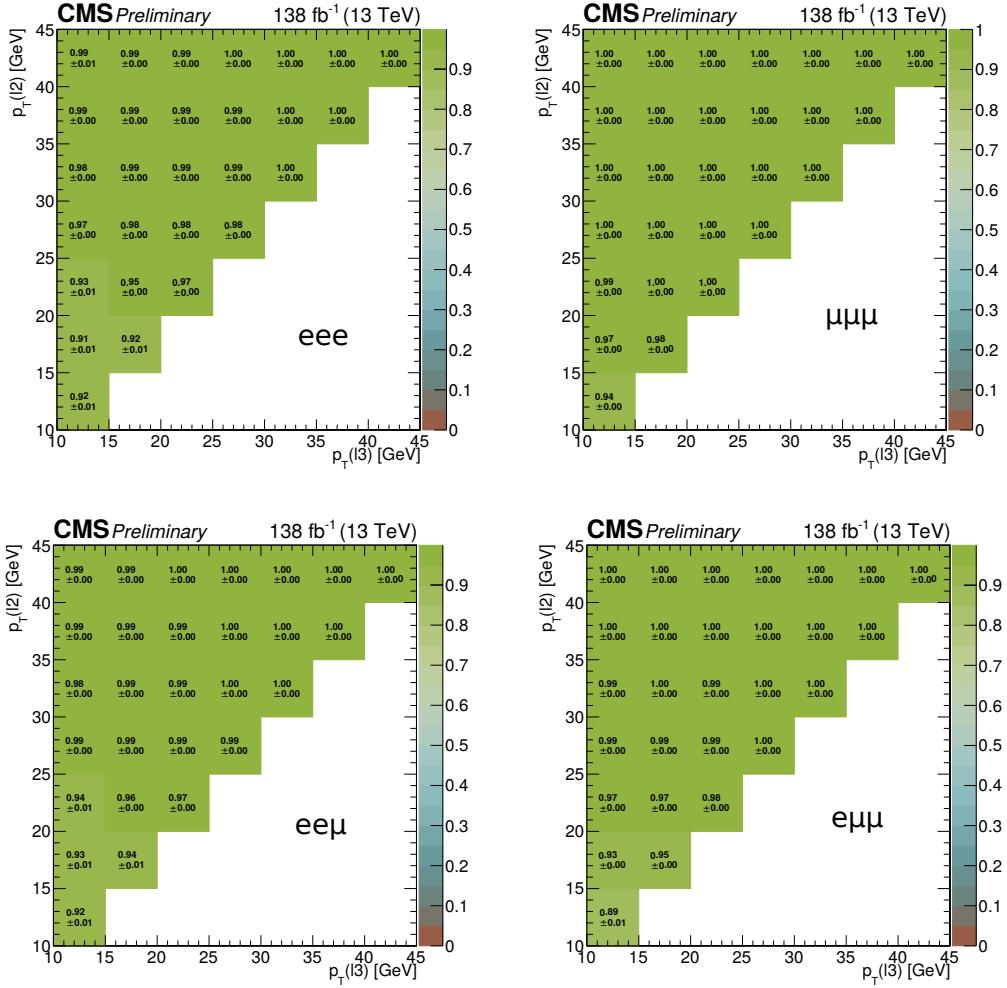


Figure 7.7: Trigger efficiency in WZ MC samples for events with 3 electrons (top left), 3 muons (top right), 1 muon and two electrons (bottom left) and 1 electron and two muons (bottom right) where the leading lepton has $p_T > 15$ GeV in Run II data.

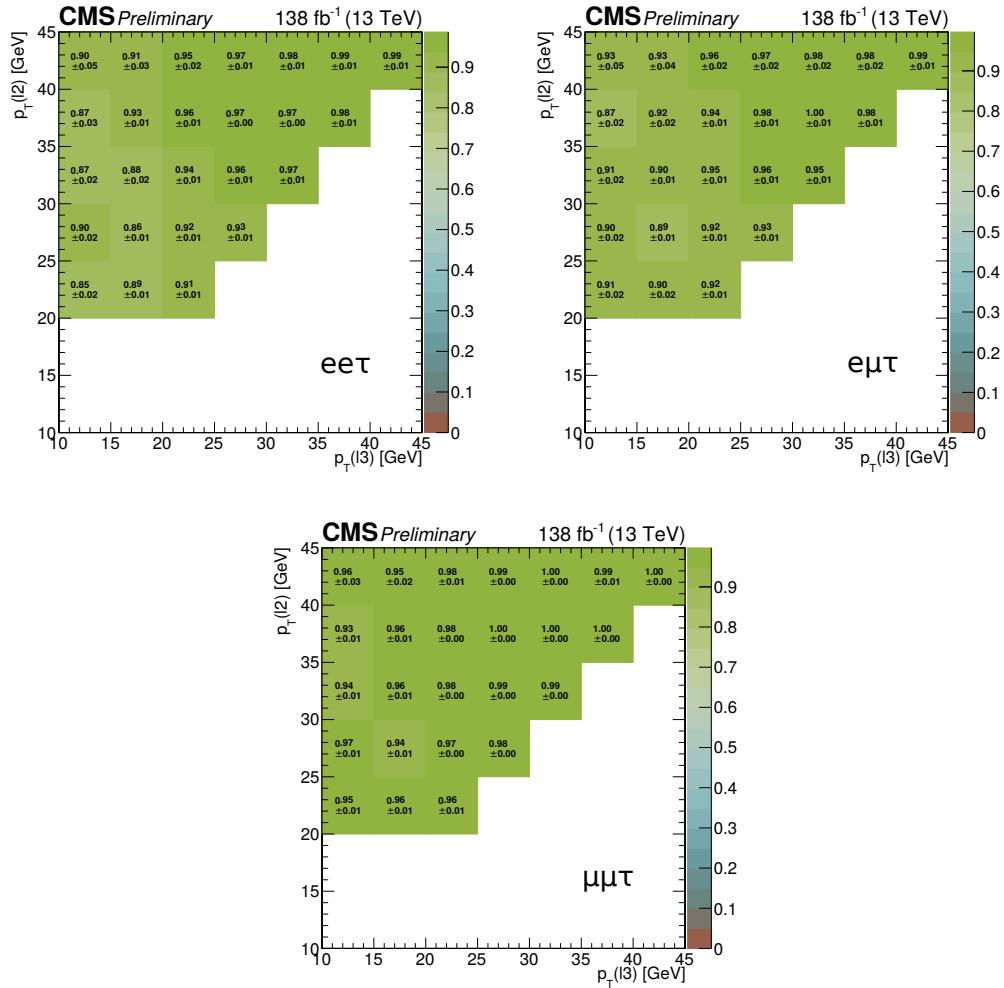


Figure 7.8: Trigger efficiency in WZ MC samples for events with 1 τ_h and 2 electrons (top left), 1 τ_h , 1 electron and 1 muon (top right), and 1 τ_h and two muons (bottom) where the leading lepton has $p_T > 15 \text{ GeV}$ in Run II data.

unrelated triggers so they are both on the same footing. One catch here is that we need to be careful not to bias our measurement due to correlations between the preselection triggers and the analysis triggers. For that reason, the trigger efficiency measurement is performed in a baseline of missing energy and jet primary datasets (MET, JetHT and MTMHT) where no lepton requirements are present and which are thus orthogonal to the data sets used in the analysis. Events are required to pass at least one of the triggers in the relevant primary datasets, an overview of which is given in Table 7.8.

Unfortunately, this comes with the disadvantage that we need to perform the measurement in a barren land full of events that have little lepton activity, making the selection of trilepton events a challenge. For that reason, all three years of data-taking as well as all final states are combined. The only splitting in final states that we will still make is in the flavour of the leading light lepton in the event. The efficiency is then once again measured by selecting events with three leptons that pass the full offline selection and checking how many pass one of the triggers. The results of this measurement in a WZ sample and the JetMET data set can

Table 7.8: List of HLT triggers used with the JetMET Data set for the measurement of the trigger efficiency in data.

| 2016 | 2017 | 2018 |
|---|---------------------------------------|--|
| HLT_MET200 | HLT_PPFJet500 | HLT_CaloMET350_HBHECleaned |
| HLT_PFMET300 | HLT_PFMET140_PFMHT140_IDTight | HLT_CaloJet500_NoJetID |
| HLT_PFMET170_HBHECleaned | HLT_PFHt500_PFMET100_PFMHT100_IDTight | HLT_AK8PFJet500 |
| HLT_PFMET120_PFMHT120_IDTight | HLT_PFHt700_PFMET85_PFMHT85_IDTight | HLT_AK8PFJet400_TrimMass30 |
| HLT_PFHt300_PFMET110 | HLT_PFHt800_PFMET75_PFMHT75_IDTight | HLT_DiJet110_35_Mjj650_PFMET110 |
| HLT_PFHt350_DiPFJetAve90_PFAalphaT0p53 | HLT_CaloJet500_NoJetID | HLT_PFHt800_PFMET75_PFMHT75_IDTight |
| HLT_PFHt400_DiPFJetAve90_PFAalphaT0p52 | HLT_AK8PFJet500 | HLT_PFHt700_PFMET85_PFMHT85_IDTight |
| HLT_PFHt400_SixJet30_DoubleBTagCSV_p056 | | HLT_PFHt500_PFMET100_PFMHT100_IDTight |
| HLT_PFHt900 | | HLT_PFHt1050 |
| HLT_PFHt650_WideJetMJJ900DEtaJJ1p5 | | HLT_PFJet500 |
| HLT_CaloJet500_NoJetID | | HLT_PFMET120_PFMHT120_IDTight |
| | | HLT_PFMET250_HBHECleaned |
| | | HLT_PFMET200_HBHE_BeamHaloCleaned |
| | | HLT_PFMETTypeOne140_PFMHT140_IDTight |
| | | HLT_PFMETTypeOne200_HBHE_BeamHaloCleaned |
| | | HLT_TripleJet110_35_35_Mjj650_PFMET110 |

be found in Figure 7.9, both for a split in leading light lepton flavour and inclusive in all final states. Good agreement is observed between data and simulation, which causes us to not apply additional corrections to the analysis but instead to apply a flat 3% uncertainty on the trigger systematic. Because the statistics in the first few bins are much lower than at high p_T , a binned statistical uncertainty is applied in the analysis on top of that 3%.

7.7 Analysis Strategy

Now that we are settled on the object selection and triggers, we can get to the heart of the analysis, namely the analysis strategy. There are multiple ways to approach signal extraction in an analysis like this. Two of these have been studied for application in this analysis and were eventually used in tandem. The first is a classic *cut-based approach*, where one selects a region enriched in the signal process, called a search region, and divides it into a plethora of bins in terms of variables that have high discriminating power between signal and background.

This helps to create bins with a lot of signal and little background as well as create bins that have the opposite effect in order to have a handle on the background in the fits. The main advantage of this strategy is that the underlying physics of the search region bins is relatively straightforward to understand. Take for example a certain reconstructed mass variable, an excess in a certain mass region could imply new physics at that mass. A disadvantage compared to the approach that will follow next is that it is very dependent on human creativity and ingenuity to cleverly design variables and search regions with the highest discriminating power possible.

The second approach, namely using *machine learning techniques*, does not suffer from this human “limitation”. One of course still has to be creative, ingenious and clever in designing the machine learning algorithm but the algorithm takes care of finding correlations between variables that lead to the highest discriminating power possible, often finding hidden connections. The main disadvantage, of course, becomes that the physics behind the machine learning algorithm becomes less straightforward to interpret.

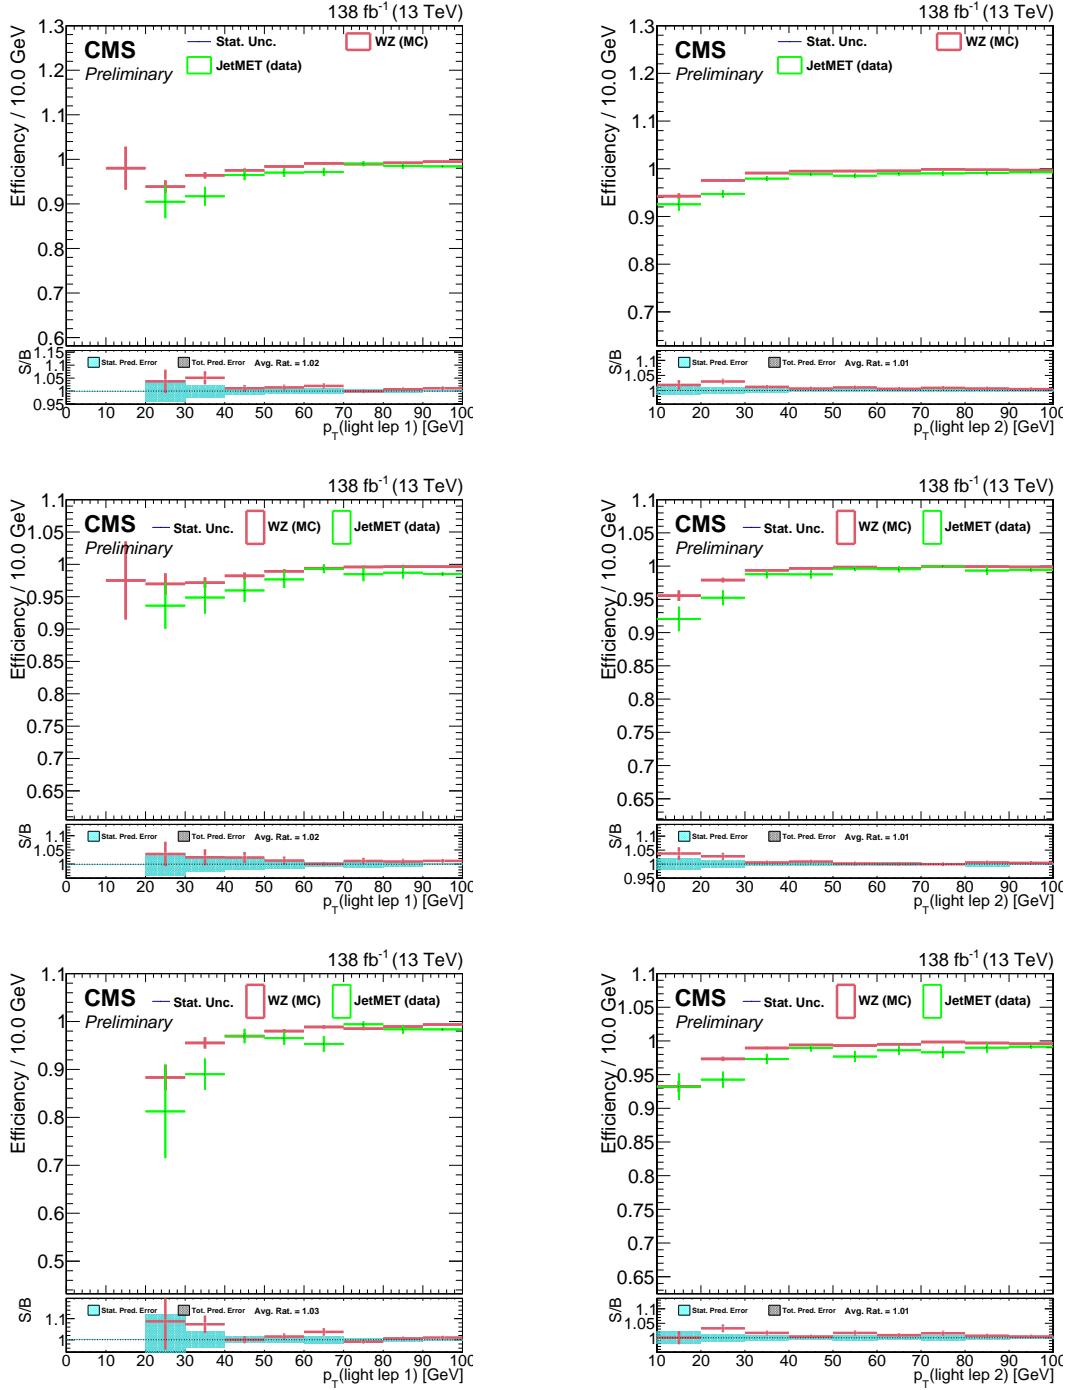


Figure 7.9: Trigger efficiency in Run II WZ MC and JetMET data samples for events with 3 leptons as a function of leading light lepton p_T (left) and subleading light lepton p_T (right). Results are shown for all final states combined (top), events where the leading light lepton is a muon (middle) and events where the leading light lepton is an electron (bottom).

Both strategies start from the same baseline search region selections that make sure our playing field will be one that is enriched in signal processes. This baseline selection is the topic of the first subsection and will promptly be followed by an overview of both approaches in the following two subsections.

7.7.1 Event selection

Processes with an HNL that has a mass lower than the W boson mass have quite different kinematics compared to models with an HNL that has a mass higher than the W mass. Therefore it is quite natural to split the analysis into two orthogonal signal regions: one to cover the first type and one to cover the latter type. They are appropriately called the *low-mass region* and the *high-mass region*. Both regions share a baseline selection before they apply their own region-specific cuts.

Baseline selection

To start off, all events should pass at least one of the triggers. Furthermore, all events that do not contain exactly three tight leptons (e , μ or τ_h), of which at most one is allowed to be a τ_h , are taken out of the signal region. In case one of the three leptons does not cross the threshold to be considered tight but does pass the FO working point, the event is not thrown away but kept in a sideband region for later use in the background estimation. If the three leptons all have the same sign, the event is discarded. Likewise, the presence of a fourth FO lepton heralds a veto of the event. Finally, the trilepton system must adhere to the offline p_T thresholds as described in Table 7.7 on top of a baseline selection of leading, subleading and trailing lepton p_T higher than 15, 10, and 10 GeV.

As our core process, as shown in Figure 7.1, does not produce jets (let alone b-jets) except for initial state radiation, a limitation on the presence of b-jets in the event is put in place. In case a jet passing our object selection with a p_T larger than 25 GeV also passes the loose b-tagging working point, the event is vetoed. Due to the overwhelming presence of SM backgrounds containing a Z boson and a limited amount of person power, the choice was made to focus on the W decay channel and place a veto on the presence of an OSSF pair that has an invariant mass within a 15 GeV window around the Z boson mass, which unfortunately also removes sensitivity to the Z decay channel of the HNL in scenarios where the HNL mass is higher than the W boson mass. Additionally, all OSSF lepton pairs in the event should have an invariant mass larger than 5 GeV. This is done to remove low-mass resonance contributions as well as photon conversions.

Low-mass region selection

Reminding ourselves of the kinematics of HNL processes with an HNL mass lower than the W boson mass as described in Section 7.3.2, we recall that the initial W will most likely be real while the W from HNL decay has to be virtual and will be on the low- p_T end of the spectrum. As the initial decay of the W boson has to distribute its energy over both the HNL and its decay product lepton, also the first lepton will have a relatively low p_T with values often found below roughly 40 GeV in the simulated samples. Therefore these processes are typically described by a generally low p_T spectrum, which shows up both in the visible lepton p_T and the p_T^{miss} , as well as a relatively low trilepton invariant mass $m(3\ell)$ with values around or below the W mass.

The three cuts employed in the low-mass region address each of these characteristics. Firstly, the leading lepton in the event should have a p_T no larger than 55 GeV. This requirement removes 71% of the $t\bar{t}$ -type events that passed the baseline selection from the background. Secondly, p_T^{miss} should reside below 75 GeV, thereby removing events with a typically larger p_T^{miss} such as WZ and $t\bar{t}$ while the signal yield is nearly untouched (< 1.5% drop in signal yield). Last but not least there is a requirement on the invariant mass of the trilepton system. This quantity should be lower than 80 GeV, which removes backgrounds with asymmetric conversions. In those cases one of the leptons from conversion runs away with most of the energy and the other lepton can be missed in reconstruction. For this type of background, the trilepton mass typically is found around the Z boson mass.

High-mass region selection

Signal processes with an HNL mass larger than the W boson mass can only be produced in the decay of a virtual W boson. While the initial lepton will suffer from this with low p_T features, there is usually more energy in the HNL decay. Therefore the leading lepton, p_T^{miss} and $m(3\ell)$ are not bound by low values and could be allowed to roam free. An orthogonal cut on at least one of them is needed, however, to keep this region separated from the low-mass region. Studies by my colleagues for the 2016-only iteration of the analysis have shown that the optimal of the three orthogonal cuts is the one on the leading lepton. Therefore a selection of $p_T(\ell_1) > 55$ GeV is enforced.

A selection on $p_T(\ell_2) > 15$ GeV and $p_T(\ell_3) > 10$ GeV is used to remove a decent portion of the background due to nonprompt leptons at the price of losing some signal acceptance. Although the $m(3\ell)$ variable is allowed to take on any value, we do keep a veto on trilepton systems that have an invariant mass within a 15 GeV of the Z boson mass in order to retain the suppression on asymmetric conversions that we also strive for in the low-mass region. Lastly, a few additional requirements are uniquely applied to events with two same-sign electrons and a muon. This is done in order to remove the large contribution of backgrounds due to charge misidentification of one of the electrons. Firstly, both electrons should pass a charge consistency cut on all methods of charge identification. Secondly, the mass of the same-sign electron pair should not be found within 15 GeV of the Z boson mass.

7.7.2 Cut-based search bins

In the most basic form, one could simply collect the expected and observed yields in a single cut-and-count bin and feed that as input in the final fit. One can however increase sensitivity to the signal by creating search region bins in variables with high discriminating power. Some bins will be more sensitive to a certain type of signal and provide boosted sensitivity in the final fit while other bins will be sensitive to a different type of signal or the background. The collection of bins can then serve as a shape input to the fit.

Just as in the event selection, the split in low-mass signal region and high-mass signal region is continued here. For both a selection of highly discriminating variables and their binning

needs to be made. While these differ for both regions, one important trend that will be found in both is a splitting into bins containing events that have an OSSF pair present and events that do not. The latter are typically characterized by very small background yields and provide large sensitivity, especially for the low-mass region, while the first contains significantly more SM background events involving Z bosons. The choices made here are based on the search region bins employed in the previous iteration of the analysis.

Low-mass region

For HNLs with a mass lower than that of the W boson, two discriminating variables form the definition of the search bins. The first of these variables is the leading lepton p_T , from which a splitting into two categories is generated by keeping events with $p_T(\ell_1)$ lower than 30 GeV separated from events with $p_T(\ell_1)$ higher than 30 GeV. The idea behind this is that when an HNL has a mass that is close to the W mass, there is very little phase space left for the lepton from the original W decay and all leptons are relatively soft. In scenarios where the HNL is so low that there is a greater mass splitting between the HNL and the W boson, the same lepton can be relatively hard. The bins with $p_T(\ell_1) < 30$ GeV will thus be more sensitive to models with low $|m_N - m_W|$ while the orthogonal bins cater to models with higher such values.

A second variable that provides very high discriminating power is the mass of the OS lepton pair that has the lowest invariant mass: $\min m(2\ell|OS)$. The power of this variable is found in the fact that the two leptons from HNL decay have OS and thus will thus cause the distribution to fall steeply at the HNL mass. The shapes of both variables in the expected background and a few signal points can be found in Figure 7.10 for three different coupling scenarios while the exact definition and names of the low-mass region search bins can be found in Table 7.9.

High-mass region

Three variables are used to define the high-mass search region bins, the choice of which is again based on the previous iteration of this analysis. There, a range of variables was conjured up and all the ones that visually showed shape differences were picked out and used to define a set of preliminary search regions. By performing a simultaneous fit to these preliminary regions and then comparing the obtained expected limits, the highest-performing variables and their combinations were selected. Out of all variables, two stood out in terms of performance. These were the $\min m(2\ell|OS)$ discussed in the low-mass region bins and the transverse mass of the remaining lepton that is not part of the minimum OS pair in combinations with the missing transverse momentum: M_T^W . It is calculated in the following way:

$$M_T^W = \sqrt{2p_T^{\text{miss}} p_T^\ell (1 - \cos \Delta\phi)} . \quad (7.6)$$

A third variable, the mass of the trilepton system $M_{3\ell}$, was found to provide additional discrimination and this mainly for the higher-end HNL mass spectrum where $M_{3\ell}$ reaches values over 1 TeV for the highest mass points considered. The shapes for these three variables

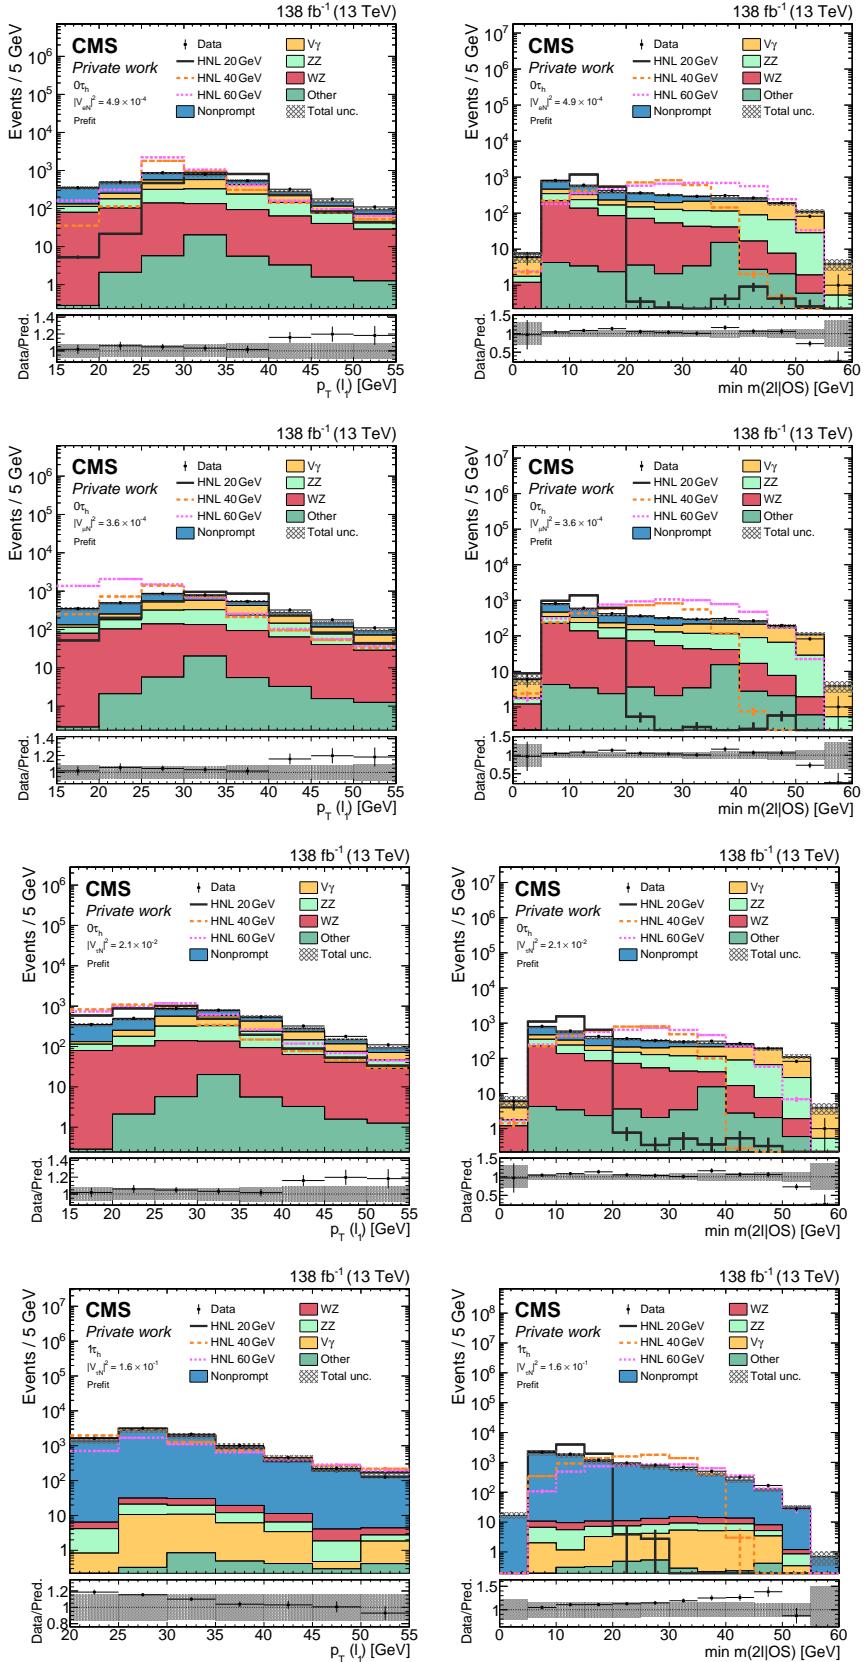


Figure 7.10: Signal and background distributions of $p_T(\ell_1)$ (left) and $\min m(2\ell|\text{OS})$ (right) in the low-mass signal region. The ratio on the bottom shows the signal yield over the background yield. Shown are the light-lepton final states with HNL models for three different masses coupling exclusively to electron (first row), muon (second row), and tau (third row) neutrinos (with identical SM backgrounds in each row), and the final states with a τ_h with HNL models for three different masses coupling exclusively to tau neutrinos (fourth row).

are shown in Figure 7.11 for three different coupling scenarios. The exact definition and naming of the search region bins from these three variables can be found in Table 7.10.

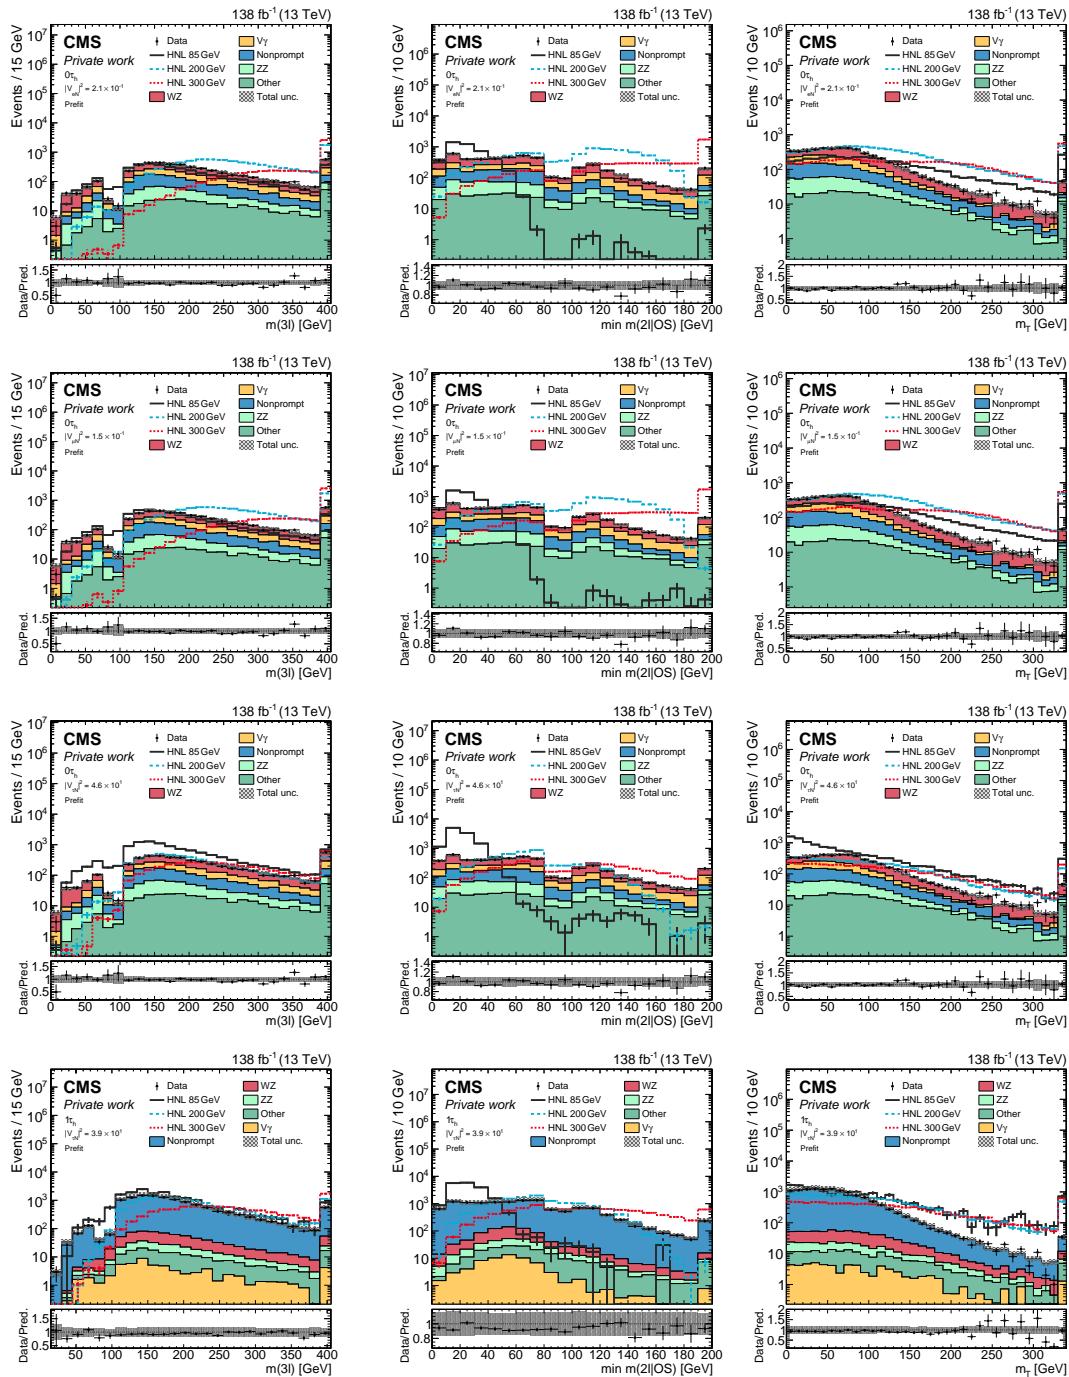


Figure 7.11: Signal and background distributions of $m(3\ell)$ (left) and $\min m(2\ell|OS)$ (middle) and M_T^W (right) in the high-mass signal region. The ratio on the bottom show the signal yield over the background yield. Shown are the light-lepton final states with HNL models for three different masses coupling exclusively to electron (first row), muon (second row), and tau (third row) neutrinos (with identical SM backgrounds in each row), and the final states with a τ_h with HNL models for three different masses coupling exclusively to tau neutrinos (fourth row).

Table 7.9: Search regions for events in the low-mass categories.

| OSSF pair | $p_T(\ell_1)$ (GeV) | $m(3\ell)$ (GeV) | $\min m(2\ell \text{OS})$ (GeV) | Bin name |
|-----------|------------------------|---------------------|------------------------------------|----------|
| No | <30 | <80 | <10 | SR La1 |
| | | | 10–20 | SR La2 |
| | | | 20–30 | SR La3 |
| | | | >30 | SR La4 |
| | 30–55 | <80 | <10 | SR La5 |
| | | | 10–20 | SR La6 |
| | | | 20–30 | SR La7 |
| | | | >30 | SR La8 |
| Yes | <30 | <80 | <10 | SR Lb1 |
| | | | 10–20 | SR Lb2 |
| | | | 20–30 | SR Lb3 |
| | | | >30 | SR Lb4 |
| | 30–55 | <80 | <10 | SR Lb5 |
| | | | 10–20 | SR Lb6 |
| | | | 20–30 | SR Lb7 |
| | | | >30 | SR Lb8 |

Table 7.10: Search regions for events in the high-mass categories.

| OSSF pair | $p_T(\ell_1)$ (GeV) | $m(3\ell)$ (GeV) | $\min m(2\ell \text{OS})$ (GeV) | M_T^W (GeV) | Bin name |
|-----------|------------------------|---------------------|------------------------------------|------------------|----------|
| No | >55 | <100 | any | <100 | SR Ha1 |
| | | | | >100 | SR Ha2 |
| | >100 | <100 | <100 | <100 | SR Ha3 |
| | | | | 100–150 | SR Ha4 |
| | | >100 | >150 | 150–250 | SR Ha5 |
| | | | | >250 | SR Ha6 |
| | | | | 100–200 | SR Ha7 |
| | >200 | <100 | any | >100 | SR Ha8 |
| | | | | | SR Ha9 |
| Yes | >55 | <75 | any | <100 | SR Hb1 |
| | | | | 100–200 | SR Hb2 |
| | | | | >200 | SR Hb3 |
| | >105 | <100 | <100 | <100 | SR Hb4 |
| | | | | 100–200 | SR Hb5 |
| | | >100 | >200 | 200–300 | SR Hb6 |
| | | | | 300–400 | SR Hb7 |
| | | >200 | >400 | >400 | SR Hb8 |
| | | | | 100–200 | SR Hb9 |
| | | | | 200–300 | SR Hb10 |
| | | | | >300 | SR Hb11 |
| | | | | | SR Hb12 |
| | >200 | <100 | <100 | <100 | SR Hb13 |
| | | | | 100–200 | SR Hb14 |
| | | >100 | >200 | 200–300 | SR Hb15 |
| | | | | >300 | SR Hb16 |

7.7.3 Boosted decision trees

The second prong of the strategy falls back on machine learning to find any exploits in the kinematic features that we could not find ourselves. Although we have skimmed over the term MVA and BDT before, we never went into detail as it was part of someone else's work. With the MVA strategy for event binning, however, the story is a bit different. In order to understand the choices behind the strategy, let us very briefly look at the basics of BDTs [201].

Introduction to boosted decision trees

Decision trees are a machine learning application that can be used for classification and regression. The concept of a decision tree has been around since the mid-20th century and the first algorithms for pattern recognition were developed in the 1980s. The idea behind it is fairly straightforward. You start from a node that imposes a certain condition on an event feature and depending on the outcome you move further to the next node connected to that outcome. This process is repeated until you hit the final node of the path that gives the final decision. In principle it could be compared to a cut-based analysis where every output node represents a category, however, in the context of this analysis we will only consider binary decision trees where an output node either tells you an event is classified as signal or as background.

It is trained by checking at each node recursively what variable provides the best separation power between signal and background events and at what splitting value this occurs. During this training, the tree is allowed to reuse previously employed variables so it can find intervals of interest instead of a single cut. While it is human-readable, quite simple in concept and is not sensitive to proper input variable manipulation, there are some drawbacks to it. It is quite a greedy algorithm that is prone to overfitting, meaning it will design a tree that works very well on the training set but focuses on a few aspects of the training set that are not general features and therefore will not be optimally trained for signal to background discrimination in new events. A change in the training set may lead to severe changes in the tree structure if the training set is not large enough, causing the algorithm to be unstable. Additionally, decision trees give discrete outputs, causing the final decision to be dependent on the purity of all previous nodes.

To overcome these drawbacks, a machine learning technique called boosting can be applied to decision trees. The idea behind boosting is to iteratively combine weak learners – i.e. classifiers with an efficiency that is not that much better than random assignment – into a single strong classifier with high efficiency. Boosted decision trees are an example where this technique works remarkably well. There are many different algorithms for creating boosted decision trees but in general, they all follow the same idea. One starts with a decision tree that is trained on a certain dataset. A weight is then assigned to the specific decision tree and the training set is modified in a way that highlights different information. The process is repeated by training a new decision tree on the modified training set and so on. In the end, one gets an ensemble of decision trees from which the final output comes as a quasi-continuous output score which is the weighted sum of the output scores of all trees in the ensemble. The main

advantages are that the collection of trees is less prone to overfitting and actually improves the predictive accuracy of the model. Because the pattern recognition is divided over a large number of decision trees, the final training will not be prone to the whims of a single tree when the training set is varied but will be more robust to noise and variations.

The most commonly used boosting algorithm in the context of boosted decision trees is called AdaBoost [202]. After a weak classifier decision tree is trained, the normalized training error is calculated using the absolute differences between the true class labels y_n and the corresponding predictions $f(x_n)$ for the input variables x_n , where n represents the index of a training event. The normalized training error is computed as $\epsilon = \sum_{n=1}^N |f(x_n) - y_n|$ where the sum runs over all training events. The decision tree weight α that is to be assigned to the tree in the weighted sum is then defined as $\alpha = \beta \times \ln \frac{1-\epsilon}{\epsilon}$ where β is a parameter that will define the learning rate. Each event in the training set is then weighted with a weight w defined as $w = w_{\text{previous}} \exp(\alpha |f(x_n) - y_n|)$ where w_{previous} is the weight from the last round of boosting. The next tree is then trained on the reweighted dataset and the process is repeated until a stopping condition is met. The final output score for a new event is then obtained as $\sum_{m=1}^M \alpha_m f_m(x)$ where m runs over all trees in the ensemble.

A second popular boosting technique is called gradient boosting and works on the principles of gradient descent. Instead of giving a higher weight to wrongly classified events as in the case of AdaBoost, gradient boosting will try to minimize a loss function in the next iteration by training for the residual $y - f(x)$ of the ensemble instead of the truth y . In this way, the new training will try to counteract the residual of the previous iteration and will gradually converge to an ensemble with accurate classification with every added tree. While gradient boosting often gives predictive accuracy that is higher than other algorithms, because it keeps trying to minimize the error with every iteration, it is prone to overfitting out the errors. At the same time, this makes them computationally a bit more demanding to train.

Boosted decision trees in HNL

How can we use these BDTs in our analysis? If you think about it, the output score of a BDT is a highly optimized discriminating variable. By defining a binning for the number of events as a function of the BDT output scores, one again has a set of search regions that are enriched in signal or background. These bins can be used as input to the simultaneous fit in the same way the search regions are used. We train different BDTs for the low- and high-mass region since they are kinematically divergent and also train a separate BDT per coupling scenario. Furthermore, because of differences in kinematics within one of the signal regions, the relevant signal masses are split up into ranges corresponding to signal scenarios where the kinematics are comparable enough. A different training is performed for each of these sub-mass ranges. Additionally, in the scenario of N - ν_τ coupling, a split is made in final states with or without a τ_h . For further bookkeeping, the notation “ $\text{BDT}(m_N, \ell, i\tau_h)$ ” will be used to denote the BDT for mass range m_N , HNL coupling to flavour ℓ and the number of τ_h in the final state.

This training is done using the TMVA package for ROOT [200]. A range of hyperparameters is defined and a training is performed for each combination. The resulting trainings are then ranked by their area under the curve (AUC)⁶ and the one that is the highest in the list and passes a visual inspection for overtraining is selected for final use. In total there are four hyperparameters that are varied:

- Boost type: To get an idea of which boost type works the best, trainings were performed using gradient boosting, AdaBoost and a variation of AdaBoost called RealAdaBoost
- Number of trees: This refers to the total number of trees to be used in the ensemble. Typically the gradient boosting algorithm can benefit from a large number of trees. However, we are dealing with limited background statistics in the training for certain regions. In such cases, a large number of trees will cause overtraining. Therefore a range starting from a very low number of trees is used in our trainings, namely the range 25–400 trees.
- Max depth: The maximum number of decision node layers in a tree. For boosted ensembles, we benefit from combining simple, weak learners. Therefore the maximum depth of the trees is kept low to the range of 2-4 layers.
- Shrinkage: This refers to the learning rate for the gradient boost algorithm. The lower this value, the lower the learning rate. A large learning rate results in large updates of the “weights” in the algorithm and thus more drastic steps. A lower learning rate typically leads to slower but more controlled convergence, translating to more robustness and less overfitting. However, lower learning rates also require a higher number of trees. For the training here, we experimented with a range of learning rates starting quite low and ending quite high, namely from 0.1 to 1.

The training was performed using MC simulation and the lack of statistics in these samples was the biggest enemy in training these BDTs. In order to battle this, simulated samples corresponding to all 3 years of data-taking are combined. The background samples consist of prompt DY, prompt $t\bar{t}$, prompt WZ and nonprompt simulated events. The last of these four backgrounds is obtained from a sideband in DY, $t\bar{t}$, and WZ simulation to once again enhance the statistics of the training sample. For each sub-mass range, the samples with relevant mass and coupling are combined. Each of the background samples is reweighted to their cross section. All signal events on the other hand are scaled to unity and normalized according to their relative sample size in order to avoid bias to a specific scenario.

As a BDT is not really sensitive to the number of input variables – it will ignore the ones that are not interesting completely – we started training with a wide array of inputs. Included in the input variables were lepton properties, reconstructed mass variables, information regarding

⁶AUC is a measure of the performance of a BDT in a binary classification task. It represents the area under the curve of the ROC plot, which plots the true positive rate against the false positive rate at different classification thresholds. AUC values range from 0 to 1, with higher values indicating better model performance. An AUC of 0.5 indicates a model with no discrimination ability, while an AUC of 1 indicates a model with perfect discrimination.

the two leading jets if present and angular information between all leptons and jets. A full overview of all variables employed can be found in Table 7.11.

Low-mass region training

In the cut-based strategy section, it was noted that there are two main types of event categories in the low-mass region: events that contain an OSSF lepton pair and events that do not. The first of these two contains the bulk of the background from e.g. DY, WZ or ZZ processes while the latter only contains background in very low quantity. This makes the events without OSSF a decidedly important contribution to the final sensitivity. As a matter of fact, the previous iteration of this analysis did not consider events with an OSSF pair in the low-mass region as the backgrounds were simply too large [99]. The reason it was added now is that with a BDT approach, it is possible to create enough separation between the large backgrounds and the signal to make these events useful again.

There are some considerations one needs to make carefully, however. If one was to merge both types of events for training, the sensitivity from events without OSSF pair gets completely washed out by the much larger background. However, when the two types of events were kept separated and different BDTs were trained on both of them, we found that the number of events in the background was simply too low in the no-OSSF type of events for any sensible training. On one hand, we do not want to wash out the sensitivity but on the other hand, the reason for this high sensitivity throws a spanner in the works in terms of BDT training. As will be seen later, the solution was to drop the BDT strategy for events without OSSF lepton pairs in favour of cut-based techniques and to move forward with the BDT training for events that do contain such a pair as they lend themselves very well to this technique.

There are two main kinematic domains for signal in the low-mass region. There are scenarios where the HNL mass is quite far removed from the W boson mass and there are scenarios where the HNL mass is relatively close to the W boson mass, leading to uncompressed and compressed scenarios respectively. To account for both types, two submass ranges are defined: one running from 10–40 GeV and one running from 50–75 GeV. These submass range definitions are the result of a process of trial and error where various combinations of mass points were trained and evaluated. The final two ranges were then selected as the having the best discriminating power and the highest expected sensitivity for the relevant mass points. Their resulting trainings are shown in Figure 7.12. The preselection employed for the training sample is the same one as described in Section 7.7.1 under the header of low-mass region selection.

High-mass region training

The situation is a bit simpler in the high-mass region. The difference in sensitivity described in the previous section is less apparent in this mass region due to the relatively higher background yields in search regions Ha when compared to the low-mass regions La. For this reason, events with and without an OSSF pair are combined here. Another difference with the low-mass

Table 7.11: Overview of all input variables for the BDTs. Variables denoted with * are only used as inputs for the low-mass region BDTs, while variables denoted with † are only used in the BDTs trained in low-mass region events with light lepton final states. The last column gives the default value used when the observable is undefined (e.g., no OSSF pair or no jets in the event).

| | Variable | Description | Default |
|---------------------|---|---|---------|
| Leptons | $p_T(\ell_i)$ | Transverse momentum of the charged leptons, with $i = 1, 2, 3$ for the leading, subleading, and trailing charged lepton | — |
| | $ \eta(\ell_i) $ | Absolute value of the pseudorapidity of the charged leptons | — |
| | $\phi(\ell_i)$ | Azimuthal angle of the charged leptons | — |
| | L_T | Sum of the p_T of all charged leptons | — |
| | $\text{flavour}(\ell) \dagger$ | flavour of each lepton, with 0 for electrons, 1 for muons, 2 for τ_h | — |
| | $\text{charge}(\ell) \dagger$ | Charge of each lepton | — |
| Masses | p_T^{miss} | Missing transverse momentum | — |
| | $\min m(2\ell \text{OS})$ | Minimum value of the invariant mass among all OS lepton pairs | — |
| | M_T^W | Transverse mass of the missing momentum and the lepton not contained in $\min m(2\ell \text{OS})$ | — |
| | $m(3\ell)$ | Invariant mass of the trilepton system | — |
| | $M_T^W(3\ell)$ | Transverse mass of the trilepton and missing momentum system | — |
| | $m(\ell_i \ell_j)$ | Invariant masses of all dilepton pairs | — |
| Jets | $\max m(2\ell \text{OSSF})$ | Maximum value of the invariant mass among all OSSF lepton pairs | -1 |
| | $p_T(j_i)$ | Transverse momentum of the two leading jets, with $i = 1, 2$ for the leading and subleading jet | 0 |
| | $ \eta(j_i) $ | Absolute value of the pseudorapidity of the jets | 0 |
| | $\phi(j_i)$ | Azimuthal angle of the jets | 0 |
| | H_T | Sum of the p_T of all jets | — |
| | N_j | Number of tight jets in the event | — |
| Angular information | $\Delta R(\ell_i, \ell_j)$ | Angular distance between all different lepton pairs | — |
| | $\Delta R[\min m(2\ell \text{OS})]$ | Angular distance between the OS lepton pair chosen for $\min m(2\ell \text{OS})$ | — |
| | $\Delta R[\max m(2\ell \text{OSSF})]$ | Angular distance between the OSSF lepton pair chosen for $\max m(2\ell \text{OSSF})$ | -1 |
| | $\Delta R(\ell_i, \vec{p}_T^{\text{miss}}) *$ | Azimuthal angle distance between each lepton and the missing transverse momentum vector | — |
| | $\Delta R(j_i, \vec{p}_T^{\text{miss}}) *$ | Azimuthal angle distance between each of the leading two jets and the missing transverse momentum vector | -1 |

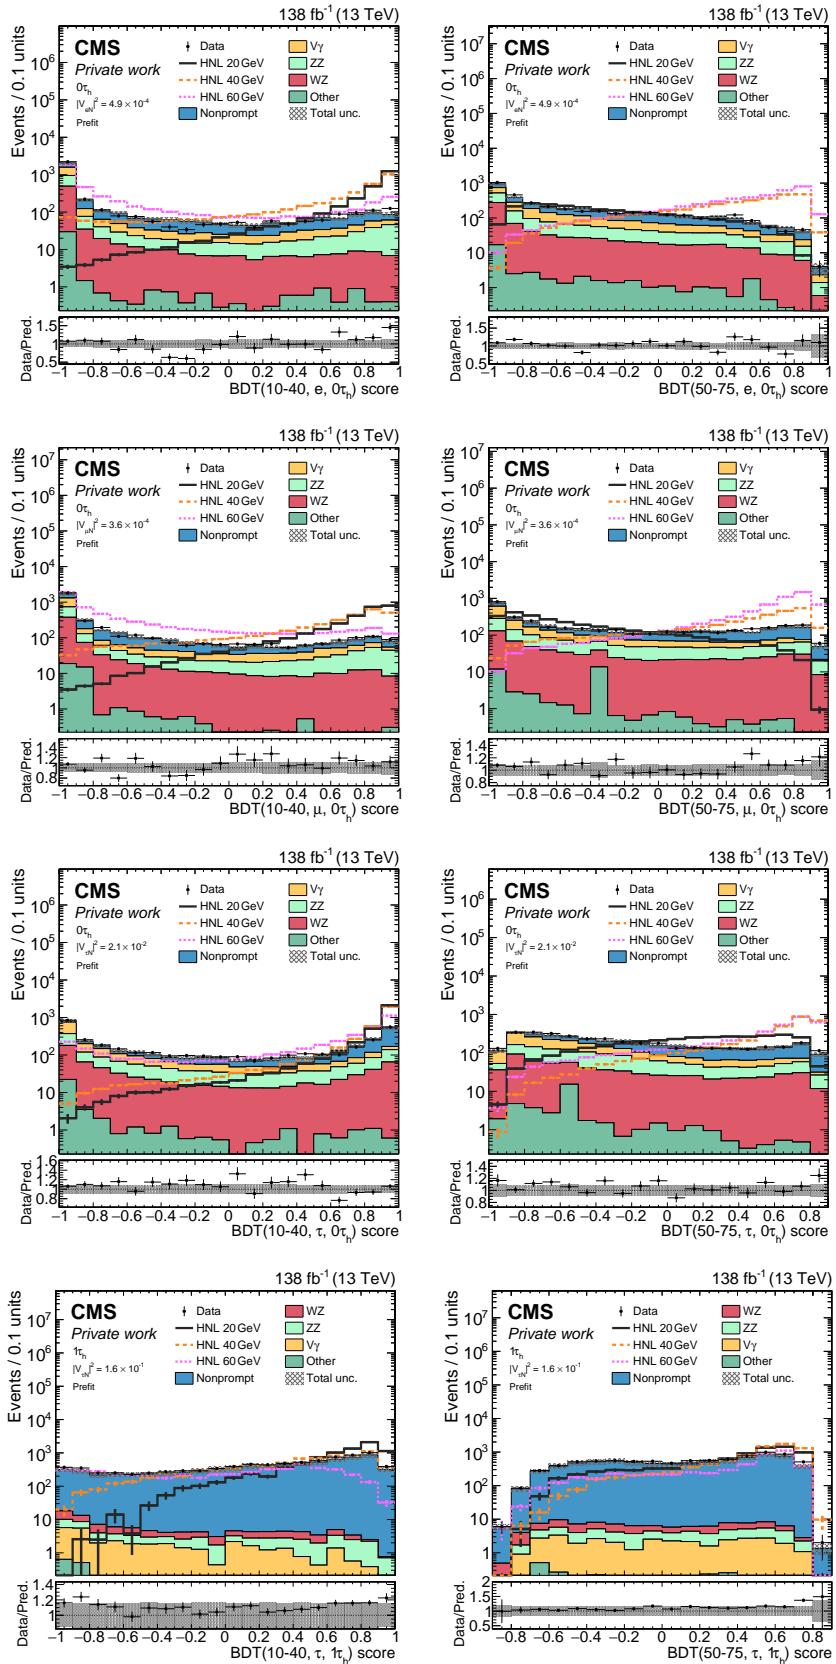


Figure 7.12: BDT output shapes for the two types of trainings in the low-mass signal region: 10–40 GeV (left) and 50–75 GeV (right). Shown per row from top to bottom as training for muon coupling, electron coupling, tau coupling in leptonic final states, and tau coupling in τ_h final states.

region is that we do not perform trainings for final states with one τ_h . The reason for this is mostly historical. We only trained preliminary tests in the initial setup of the framework. As we noticed the trainings were particularly tricky and we could not get good results in our tests, we dropped them as we thought these final states would no longer be part of our analysis anyway for collaboration reasons. We added the τ_h final states in the high-mass region again afterwards when we saw the increase in sensitivity they brought but had no time to optimize a BDT for these final states.

Also here a few m_N ranges are defined for training, namely 85–100 GeV, 100–250 GeV and 250–400 GeV. No trainings were performed for masses above 400 GeV as we found there the BDT strategy was outperformed by the cut-based strategy. The reason for this is that above this mass point, the cross section of the signal becomes too small for the signal to overcome the background in the tail of the BDT score. The resulting trainings can be found in Figure 7.13. Once again the preselection for the training sample followed the ones as defined in Section 7.7.1 under the header of high-mass region selection.

7.7.4 Combined strategy

The final sensitivity on the cross section for each HNL mass point will be obtained by performing a simultaneous fit of the signal strength to a selection of distributions. These distributions are selected from the two abovementioned strategies per fit based on the optimal expected limits for each set of distributions. An overview of the distributions going into the fit for each HNL scenario can be found in Table 7.12.

For scenarios where the HNL mass is found below the W boson mass, input distributions from the cut-based strategy and the BDT strategy are used simultaneously. Events without an OSSF pair (which are contained in search regions La) provide high sensitivity but did not lend well to the BDT strategy. Events with an OSSF pair, however (contained in search regions Lb), suffer from large backgrounds but gain big improvements from the BDT technique. Therefore we select the distribution of search region bins La for events with no OSSF and the BDT relevant to the signal scenario for the other type of events and use both shapes simultaneously in the fit.

For scenarios in this mass range with exclusive coupling to a light lepton flavour, we only regard these distributions for events that have a purely light lepton final state. Events with eee and $ee\mu$ final states are kept separate from events with $\mu\mu\mu$ and $e\mu\mu$ final states in the fit. In the end, a simultaneous fit that uses 4 input distributions (BDT score and search regions for both final state types) is performed. The idea behind these final state splittings is that in models where there is only coupling to e.g. electron neutrinos, the first two leptons in the event can only be electrons. These models are sensitive to the $eee/ee\mu$ final states and insensitive to the $\mu\mu\mu/e\mu\mu$ final states. Both final states are kept into the final fit however so that there is a good handle on background-dominated bins from the latter final states.

For scenarios where the HNL mass is lower than the W boson mass but where there is exclusive coupling to τ -neutrinos, these light lepton final states are topped up with the relevant

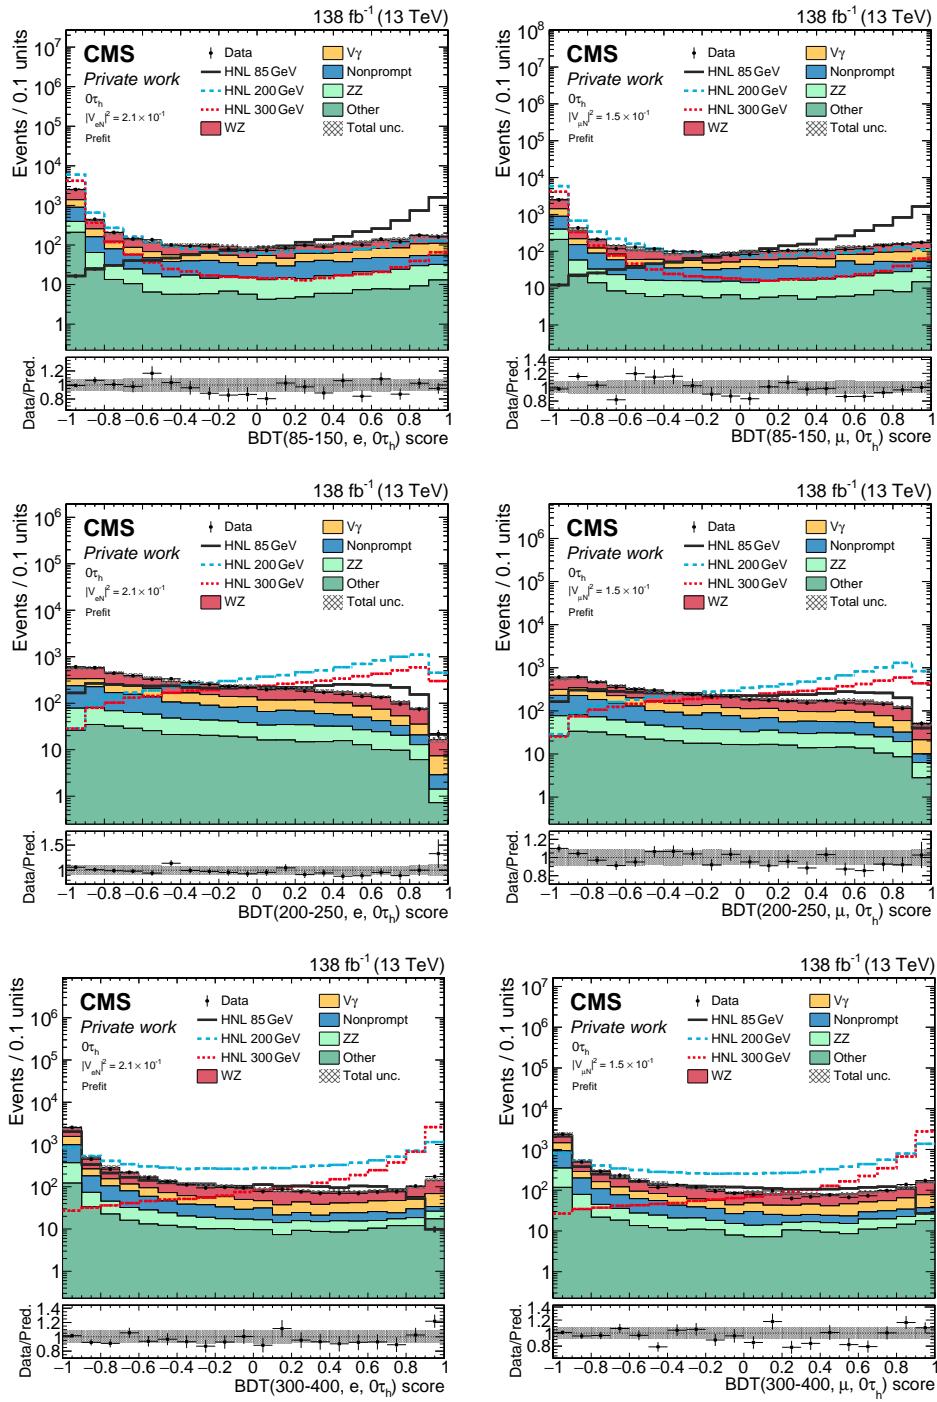


Figure 7.13: BDT output shapes for the three types of trainings in the high-mass signal region: 85–100 GeV (top), 100–250 GeV (middle) and 250–400 GeV (bottom). Shown per column from left to right as training for muon coupling, electron coupling.

search region and BDT output score distributions as described in the last paragraph for $ee\tau_h$ events, $\mu\mu\tau_h$ events and $e\mu\tau_h$ events separately. The last category of events is further split into events with OS electron and muon and those with same-sign electron and muon. Again all these distributions are combined into a single shape.

For signal scenarios with $m_W < m_N < 400$ GeV and exclusive coupling to light lepton neutrinos, the BDT strategy vastly outperformed the search bin strategy. Therefore a simultaneous fit is performed to the relevant BDT output score for that specific scenario, once again split into $eee/ee\mu$ and $\mu\mu\mu/e\mu\mu$ final states, and the search regions were not regarded.

For signal scenarios with $m_N > 400$ GeV and exclusive coupling to light lepton neutrinos, the BDT strategy starts to suffer from the tiny cross sections of the signal at these high masses. The signal peak in the tail gets lost in the background that is still making its way there. It turns out that the cut-based strategy is able to provide better sensitivity for these masses. This is the reason that the input distribution to the simultaneous fit in these scenarios is made up of search region bins Ha1–9 and Hb1–16. Also here, only light lepton final states are used and they are split up into $eee/ee\mu$ and $\mu\mu\mu/e\mu\mu$ final states.

Last but not least there are the scenarios with $m_N > m_W$ and exclusive coupling to tau-neutrinos. For these signal types, we do not use the BDT strategy. The input distributions to the final fit consist of search region bins Ha1–9 and Hb1–16 for events split into $eee/ee\mu$, $\mu\mu\mu/e\mu\mu$, $ee\tau_h$, $ee\mu$ and $e\mu\tau_h$ final states where the last category again is split into same-sign and opposite-sign electron and muon. The final shape is a combination of the search region shapes of all these different final states appended into a single shape.

7.8 Background Estimation

At this point, we have a good handle on object reconstruction, event selection and signal extraction but we are missing a handle on the background prediction. Let us delve a bit deeper into that aspect of the analysis in this section. As one might have guessed, this search deals with similar backgrounds as the SUSY search in Chapter 6. A large portion of the expected background yield comes from events with one or more nonprompt leptons, which constitute the largest background in the low-mass regions as well as the single τ_h channels in the high-mass regions. Other dominating backgrounds are WZ , ZZ where one lepton is not reconstructed and asymmetric photon conversion events from i.e. $Z\gamma$. They will be estimated from simulation and validated in control regions. Other minor backgrounds such as the production of $t\bar{t}$ in association with SM bosons or processes involving three SM bosons are also estimated from simulation but we do not reconstruct control regions for them.

7.8.1 Nonprompt background

Once again we opt for the data-driven “fake rate” technique as described in Section 6.7.1 to estimate the nonprompt backgrounds in the analysis. Overall the methodology applied is very

Table 7.12: Summary of the selections, categories, and distributions used in the template fit for each of the HNL signal points.

| HNL model | Selection | Categories | OSSF | Fitted distributions |
|------------------------|-----------|------------|------|----------------------------------|
| (10–40 GeV, e) | low-mass | $0\tau_h$ | no | La1–8 |
| | | $0\tau_h$ | yes | BDT(10–40, e, $0\tau_h$) |
| (10–40 GeV, μ) | low-mass | $0\tau_h$ | no | La1–8 |
| | | $0\tau_h$ | yes | BDT(10–40, μ , $0\tau_h$) |
| (10–40 GeV, τ) | low-mass | $0\tau_h$ | no | La1–8 |
| | | $0\tau_h$ | yes | BDT(10–40, τ , $0\tau_h$) |
| | | $1\tau_h$ | no | La1–8 |
| | | $1\tau_h$ | yes | BDT(10–40, τ , $1\tau_h$) |
| (50–75 GeV, e) | low-mass | $0\tau_h$ | no | La1–8 |
| | | $0\tau_h$ | yes | BDT(50–75, e, $0\tau_h$) |
| (50–75 GeV, μ) | low-mass | $0\tau_h$ | no | La1–8 |
| | | $0\tau_h$ | yes | BDT(50–75, μ , $0\tau_h$) |
| (50–75 GeV, τ) | low-mass | $0\tau_h$ | no | La1–8 |
| | | $0\tau_h$ | yes | BDT(50–75, τ , $0\tau_h$) |
| | | $1\tau_h$ | no | La1–8 |
| | | $1\tau_h$ | yes | BDT(50–75, τ , $1\tau_h$) |
| (85–150 GeV, e) | high-mass | $0\tau_h$ | any | BDT(85–150, e, $0\tau_h$) |
| (85–150 GeV, μ) | high-mass | $0\tau_h$ | any | BDT(85–150, μ , $0\tau_h$) |
| (85–150 GeV, τ) | high-mass | $0\tau_h$ | any | Ha1–9, Hb1–16 |
| (85–150 GeV, τ) | high-mass | $1\tau_h$ | any | Ha1–9, Hb1–16 |
| (200–250 GeV, e) | high-mass | $0\tau_h$ | any | BDT(200–250, e, $0\tau_h$) |
| (200–250 GeV, μ) | high-mass | $0\tau_h$ | any | BDT(200–250, μ , $0\tau_h$) |
| (200–250 GeV, τ) | high-mass | $0\tau_h$ | any | Ha1–9, Hb1–16 |
| (200–250 GeV, τ) | high-mass | $1\tau_h$ | any | Ha1–9, Hb1–16 |
| (300–400 GeV, e) | high-mass | $0\tau_h$ | any | BDT(300–400, e, $0\tau_h$) |
| (300–400 GeV, μ) | high-mass | $0\tau_h$ | any | BDT(300–400, μ , $0\tau_h$) |
| (300–400 GeV, τ) | high-mass | $0\tau_h$ | any | Ha1–9, Hb1–16 |
| (300–400 GeV, τ) | high-mass | $1\tau_h$ | any | Ha1–9, Hb1–16 |
| (>400 GeV, e) | high-mass | $0\tau_h$ | any | Ha1–9, Hb1–16 |
| (>400 GeV, μ) | high-mass | $0\tau_h$ | any | Ha1–9, Hb1–16 |
| (>400 GeV, τ) | high-mass | $0\tau_h$ | any | Ha1–9, Hb1–16 |
| (>400 GeV, τ) | high-mass | $1\tau_h$ | any | Ha1–9, Hb1–16 |

similar to the SUSY analysis. We will briefly go over the results and highlight any differences in obtaining them between this analysis and the SUSY search.

Background from nonprompt light leptons

The tight-to-loose ratio for nonprompt light leptons is again measured separately for electrons and muons in a QCD-enriched region characterized by the presence of one light lepton passing the FO and at least one jet that is spatially separated from it by $\Delta R > 0.7$. There are no other loose leptons allowed in this selection. In similar fashion as the search from the previous chapter, differences in light-flavour and heavy-flavour jet fakes are taken into account by using a FO working tuned to the lepton MVA selection and for isolation effects by measuring the fake rate as a function of the p_T of the mother parton from which the jet originated. The same type of study of the discontinuous jump in the average p_T spectrum as a function of the lepton MVA score due to the use of uncorrected p_T for light leptons that pass the tight working point and parton p_T for those that do not as described in Section 6.7.1 was performed. From this, the following cone correction factors were found:

$$\begin{aligned} p_T^{\text{parton}} &= 0.66 \cdot p_T / p_T^{\text{ratio}} && \text{for muons,} \\ p_T^{\text{parton}} &= 0.72 \cdot p_T / p_T^{\text{ratio}} && \text{for electrons.} \end{aligned} \quad (7.7)$$

In order to validate the method using simulation truth information, fake rate maps are first measured in simulation in the QCD region described before. An overview of the resulting fake rate maps in 2017 simulation can be found in Figures 7.14 and 7.15. Similar maps were found for the other years of data-taking. The next step is to perform closure tests in a very general selection using these fake rates in a simulated DY and a simulated $t\bar{t}$ sample. In both samples, we require the presence of at least two FO leptons which need to have same-sign if they are the only two leptons in the event. At least one of them should be nonprompt, which is decided using simulation truth by checking for the presence of generator-level light leptons in the vicinity of our candidates. The resulting closure tests in 2017 simulation can be found in Figure 7.16. Here the nonprompt leptons can be either electrons or muons. Similar behaviour was observed in the other years of data-taking.

Background from nonprompt taus

In the previous chapter, it was described how the contribution from nonprompt taus was estimated from data using two sets of fake rate maps, one measured in a DY control region and one in a $t\bar{t}$ control region. The final fake rate applied in a certain region depended on the dominating nonprompt background type. Most of the time these were light-flavour jet nonprompts for OSSF events and heavy-flavour jet nonprompts for different flavour (DF) events. We found that in this analysis it was not always clear during development how to make this split. While the principle behind nonprompt tau prediction stayed the same – measure a fake rate in DY and $t\bar{t}$ and apply accordingly – application evolved into the usage of a weighted sum of fake rate weights. The DY and $t\bar{t}$ fake rates were summed with a weight that corresponds to the estimated fraction for each type of nonprompt in simulation in the

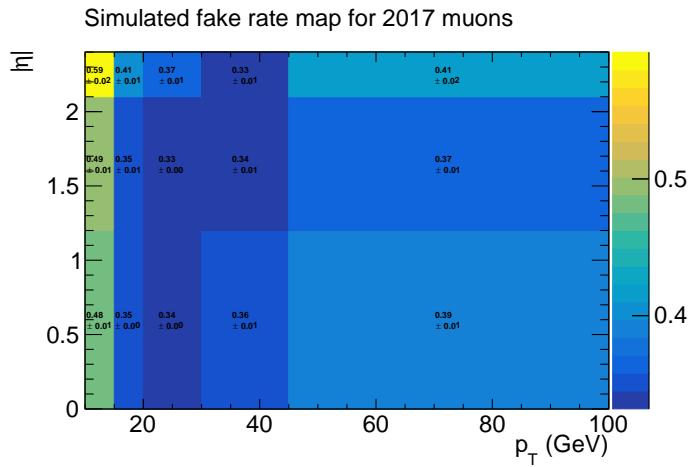


Figure 7.14: Fake rate maps obtained from 2017 QCD multijet simulation for muons. Courtesy of Luka Lambrecht.

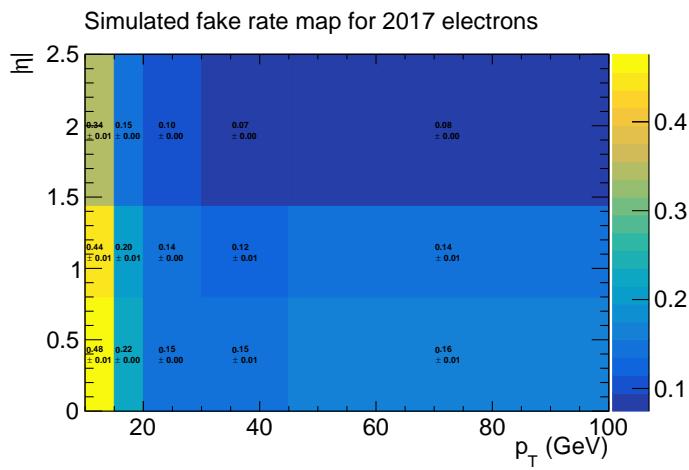


Figure 7.15: Fake rate maps obtained from 2017 QCD multijet simulation for electrons. Courtesy of Luka Lambrecht.

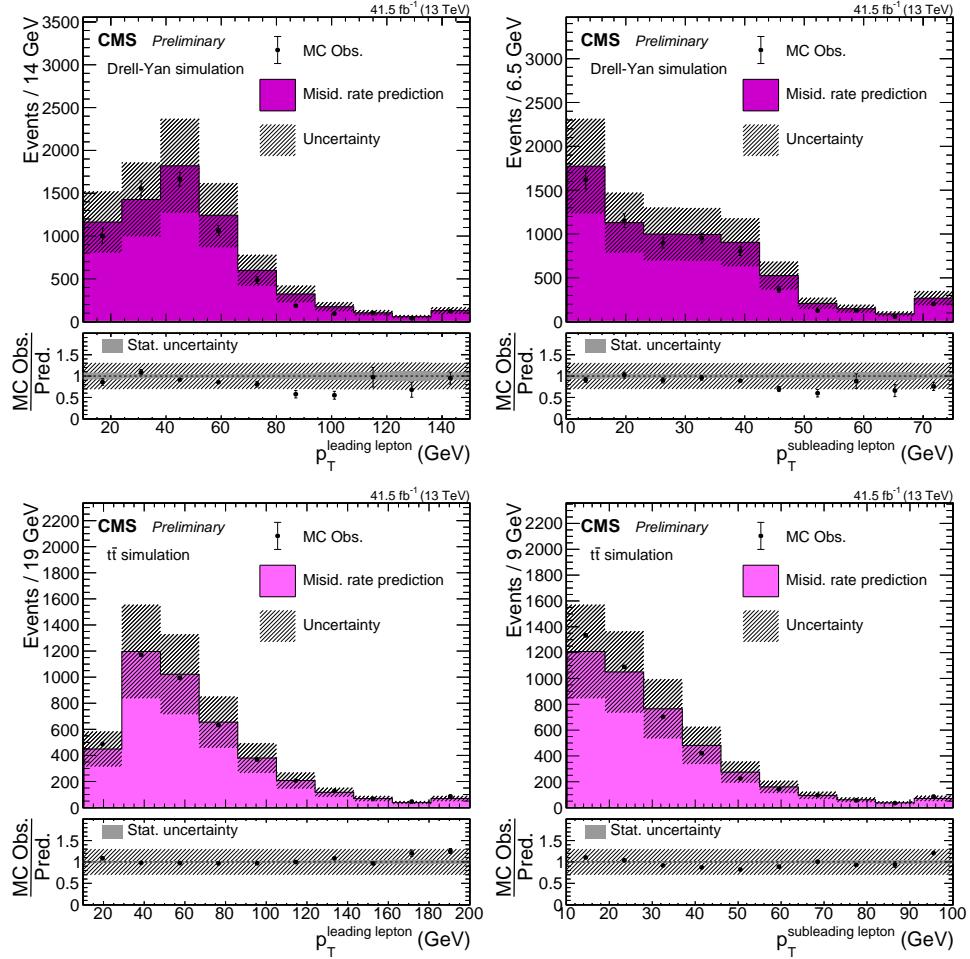


Figure 7.16: Simulated closure test result for electrons and muons combined in DY (top) and $t\bar{t}$ (bottom) simulated events according to the 2017 data-taking. Shown are the p_T distributions of the leading (first and third column) and subleading (second and fourth column) lepton in the event. Courtesy of Luka Lambrecht.

signal region. The relative contributions for each type of nonprompt τ_h is summarized in Table 7.13 per year. It can be seen that the overwhelming majority of the fakes originate from DY-type fakes. In the high-mass region, however, the contribution of $t\bar{t}$ -type fakes is not negligible. Large differences in the latter region can occur depending on the final state, with relative contributions of up to 60% by $t\bar{t}$ -type nonprompt τ_h events in final states containing an electron, a muon and a τ_h lepton while final states with two light leptons of the same flavour were dominated by DY-type nonprompts. For this reason, the final weights are made final state-dependent.

The measurement regions for the fake rate maps both in simulation and data are more or less the same as in the last chapter. The **DY measurement region** consists of a selection of two OSSF light leptons that have an invariant mass within a 15 GeV window around the Z mass along with a FO τ_h . The presence of any other FO lepton warrants a veto of the event. Additionally, the missing transverse momentum in the event should be below 50 GeV. This last cut is in place to reduce the contribution of prompt events. It was found that this region is 99.6% pure in DY events. The **ttbar measurement region** on the other hand requires the presence of a tight light lepton pair of opposite flavour and opposite sign along with a FO τ_h . The invariant mass of the light lepton pair should be above 20 GeV in order to suppress QCD contributions. At least one tight b-jet should be present. With all selections in place, we see that the relative contribution of $t\bar{t}X$ to this region is 96.1%, followed by tX with 3.6%. The resulting fake rate maps in both data and simulation for the 2017 dataset can be found in Figure 7.17. The maps measured in the 2016 and 2018 datasets contain comparable values.

The validation of these fake rates in both simulation – in order to test the validity of the method – and in data – in order to test the validity of the fake rate maps – is done in three different control regions. The first of these is a **general closure region** that is designed for validation in simulation and attempts to be as broad as possible in phase space. It is characterized by the presence of three FO objects, at of which at least one is required to be a nonprompt τ_h at the generator level. What this last statement requires is for the τ_h to not be matched to

Table 7.13: Relative contributions of nonprompt τ events from DY-type and $t\bar{t}$ -type events in the signal regions.

| | | Fraction DY | Fraction $t\bar{t}$ |
|--------------|-----------------|-------------|---------------------|
| low-mass SR | 2016 (pre-VFP) | 0.987 | 0.013 |
| | 2016 (post-VFP) | 0.980 | 0.02 |
| | 2017 | 0.982 | 0.018 |
| | 2018 | 0.982 | 0.018 |
| high-mass SR | 2016 (pre-VFP) | 0.859 | 0.141 |
| | 2016 (post-VFP) | 0.826 | 0.174 |
| | 2017 | 0.798 | 0.202 |
| | 2018 | 0.783 | 0.217 |

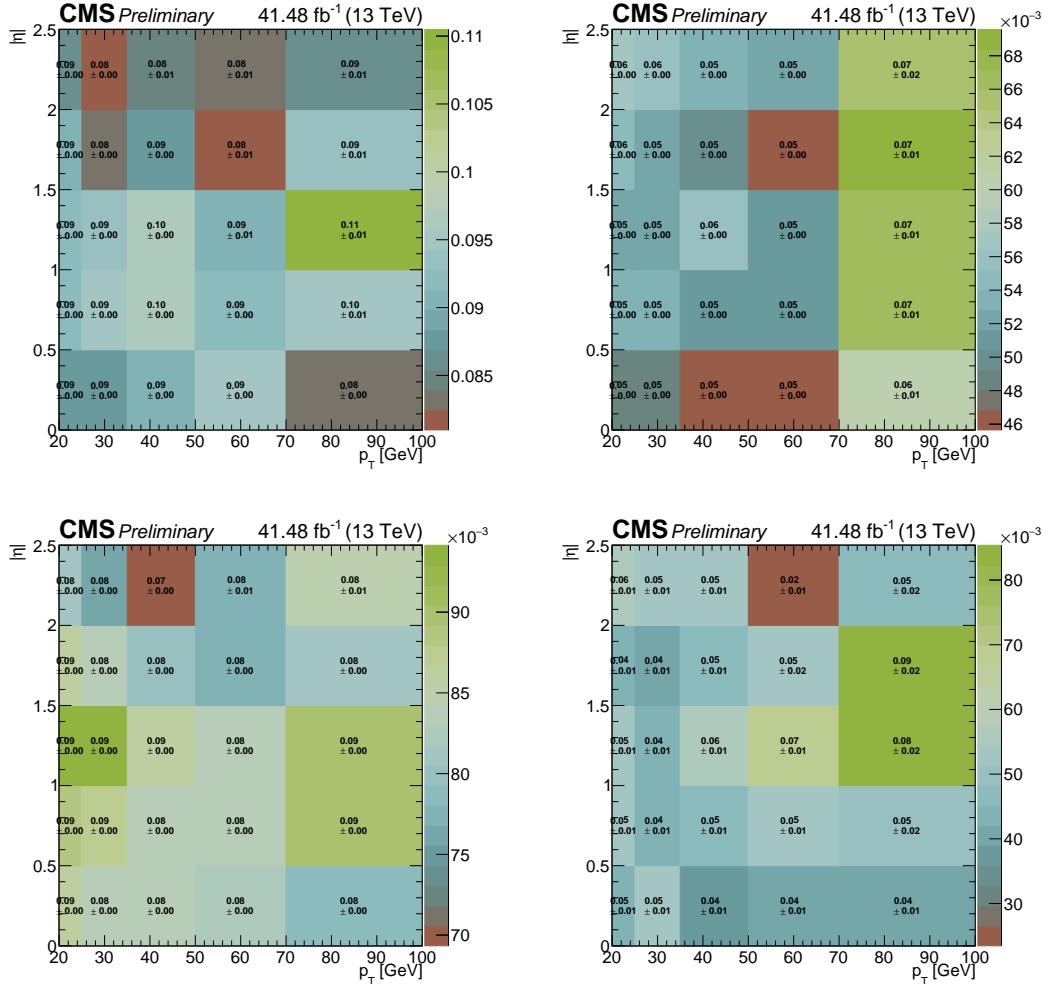


Figure 7.17: Tight-to-loose ratio for nonprompt hadronically decayed tau leptons as measured in the DY (left) and $t\bar{t}$ (right) control region for the 2017 dataset. The upper row shows the results in 2017 simulation samples, while the lower row shows the results in 2017 data samples.

electrons, muons or genuine τ 's in simulation. Furthermore, there should be no b-jets in the event and the leptons should pass the offline requirements as defined in the trigger section. The closure tests for the fake rate maps as measured in DY are evaluated in a DY MC sample with the aforementioned selection applied. Likewise, the validation for the $t\bar{t}$ fake rate maps is performed in $t\bar{t}$ simulation in this general selection. The results can be found for 2017 samples in Figures 7.18 and 7.19. Closure tests in the other years show similar results. As usual, the observed contribution represents the number of events where all leptons are tight and the predicted distribution shows the prediction from the sideband events with the proper fake rate weight value applied.

The next two control regions aim to validate both types of fake rate maps in data. The first of these is called the **DY closure region**. It is defined in almost the exact same way as the DY measurement region with the exception of the p_T^{miss} cut, which has been inverted, and the presence of a loose b-jet veto. The purity of this region in DY events is found to be at 98.6%. The closure tests in data using the DY fake rate are shown in Figure 7.20. The second region is called the **$t\bar{t}$ closure region**. After playing around with a few variables, it turned out to be quite hard to create a control region that was both orthogonal to the signal region and the $t\bar{t}$ measurement region. Since the fraction of light jet fakes in the signal region is so much higher than that of the heavy jet fakes and since there is good agreement in the closure tests in simulation, we have opted to keep the $t\bar{t}$ closure region the same as the measurement region and validate for the presence of any systematic shifts in distributions not relevant to the fake rate map binning. The resulting agreement can be found in Figure 7.21. In both the DY and $t\bar{t}$ closure tests, the predicted distribution is obtained from a sideband in data where the τ_h in the event passes the FO selection but not the tight selection and the prompt contribution is obtained from simulation. For all closure test types, a good agreement between observed and predicted distributions is observed in all four datasets.

7.8.2 WZ background

The production of a W boson along with a Z boson where both of them decay leptonically does not look unlike an HNL event due to the presence of three charged leptons and missing energy from the neutrino in the W decay. It is an important background to the analysis, especially in the high-mass region, and it is important that we understand and validate our prediction. The expected yield from WZ events in both signal regions is obtained from simulation. In order to validate how well our WZ simulation predicts reality, we define a control region that is enriched in WZ events and is defined as:

- exactly three tight light leptons with $p_T > 25, 15, 10 \text{ GeV}$;
- no fourth FO light lepton;
- two leptons have OSSF and $|m(2\ell|\text{OSSF}) - m_Z| < 15 \text{ GeV}$;
- no loose b-jets;
- $p_T^{\text{miss}} > 50 \text{ GeV}$ (to account for neutrino from W decay);

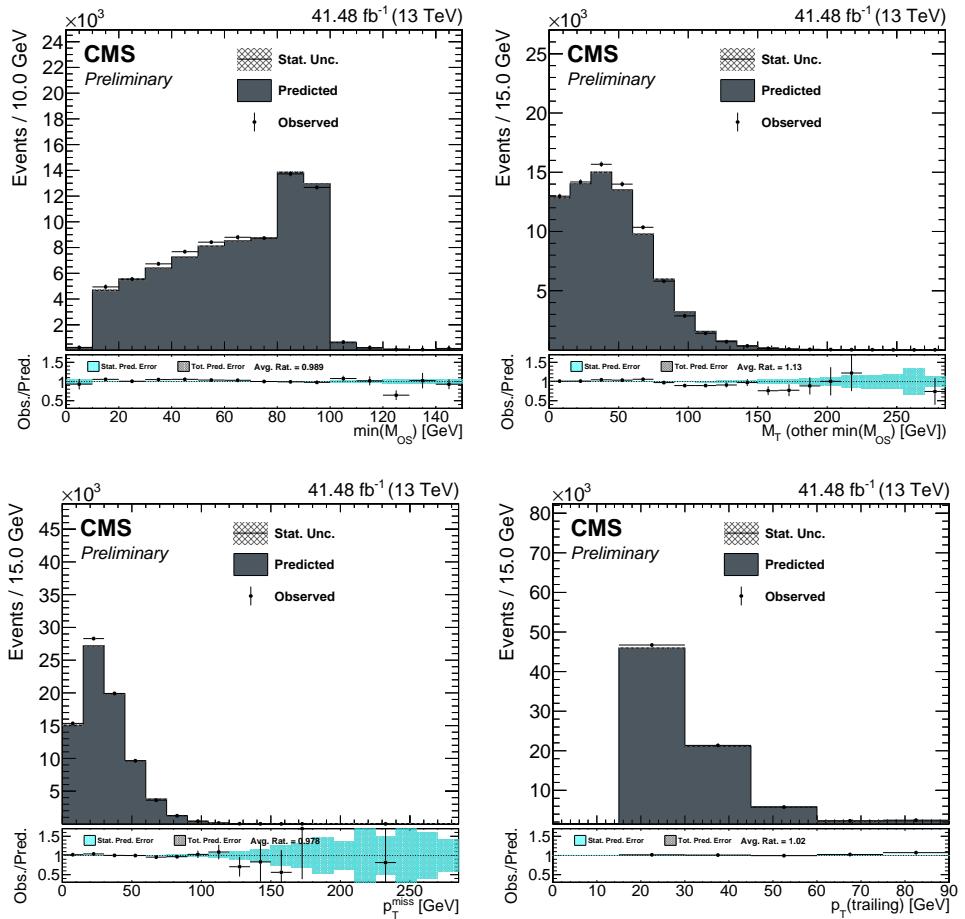


Figure 7.18: Closure tests of the tau fake rate method using a fake rate measured in DY in simulated DY samples corresponding to the 2017 dataset. Shown as a function of $\min m(2\ell|OS)$ (top left), M_T^W (top right), p_T^{miss} (bottom left) and transverse momentum of the trailing lepton (bottom right). The event selection consists of events passing a general trilepton selection.

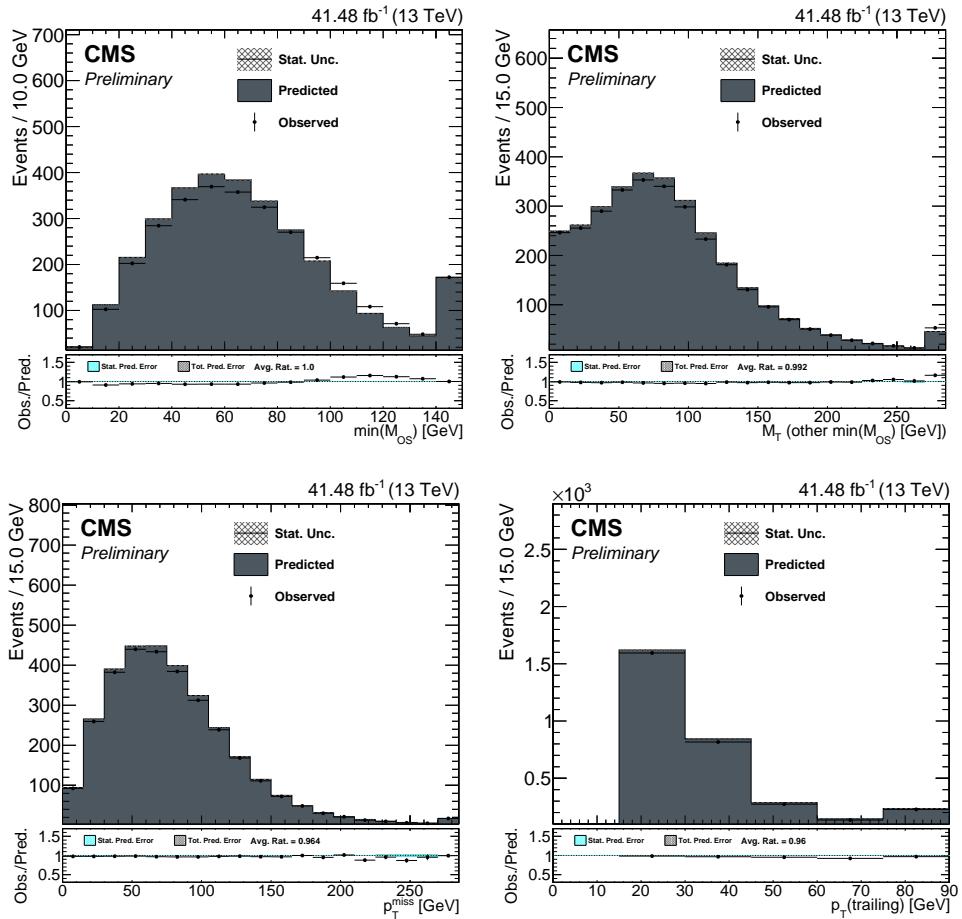


Figure 7.19: Closure tests of the tau fake rate method using a fake rate measured in $t\bar{t}$ simulated $t\bar{t}$ samples corresponding to the 2017 dataset. Shown as a function of $\min m(2\ell|OS)$ (top left), M_T^W (top right), p_T^{miss} (bottom left) and transverse momentum of the trailing lepton (bottom right). The event selection consists of events passing a general trilepton selection.

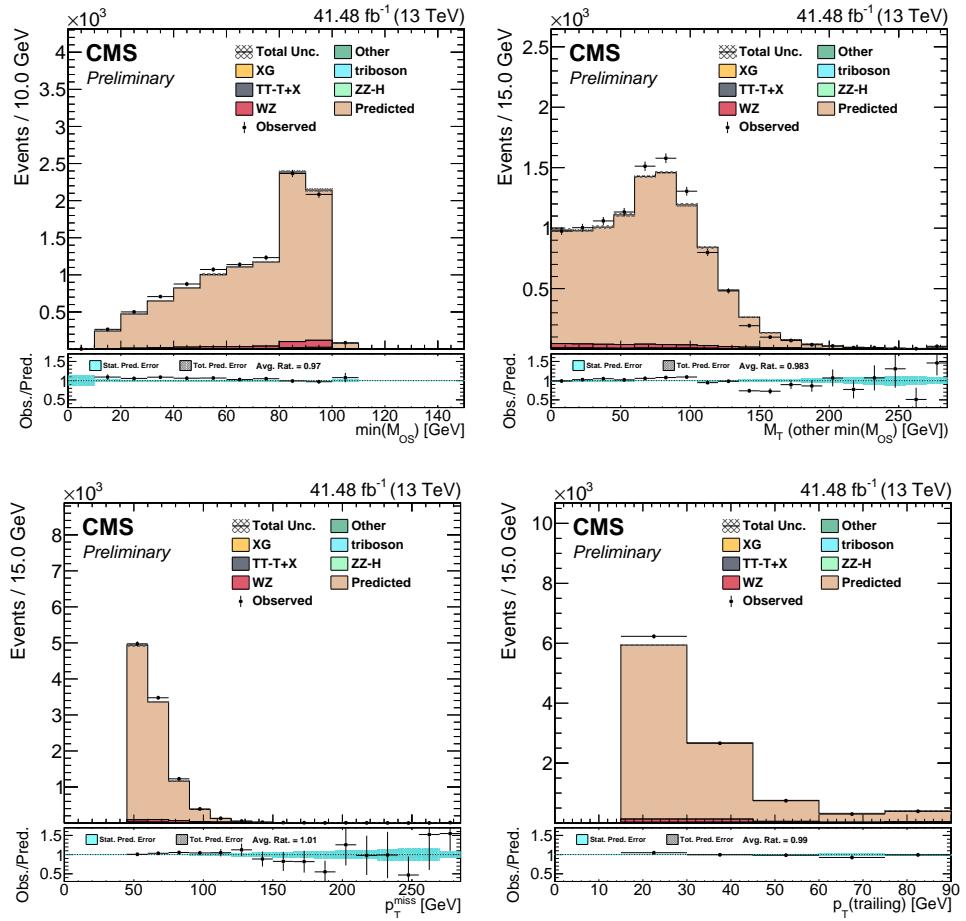


Figure 7.20: Closure tests of the tau fake rate method in the DY closure region using a fake rate measured in DY in 2017 data. Shown as a function of $\min m(2\ell|OS)$ (top left), M_T^W (top right), p_T^{miss} (bottom left) and transverse momentum of the trailing lepton (bottom right).

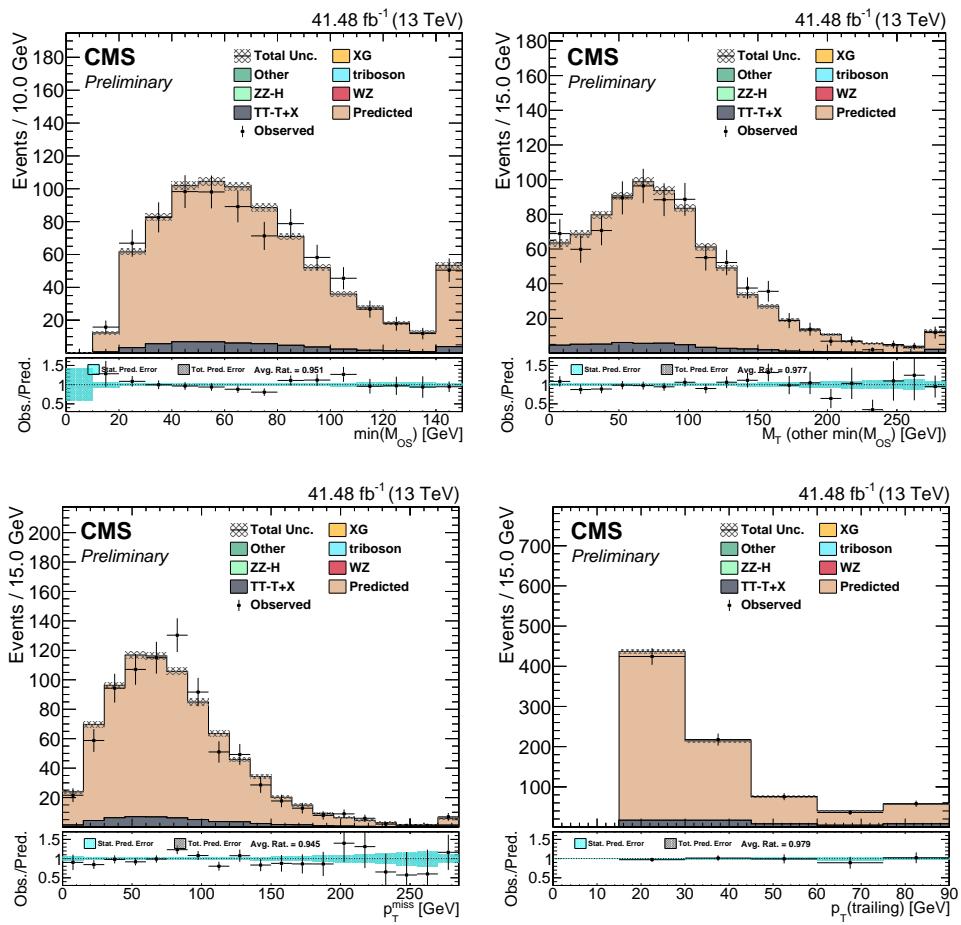


Figure 7.21: Closure tests of the tau fake rate method in the $t\bar{t}$ closure region using a fake rate measured in $t\bar{t}$ in 2017 data. Shown as a function of $\min m(2\ell|OS)$ (top left), M_T^W (top right), p_T^{miss} (bottom left) and transverse momentum of the trailing lepton (bottom right).

- trilepton invariant mass $|m(3\ell) - m_Z| > 15 \text{ GeV}$ (to remove $Z\gamma$).

This selection is about 80% pure in events that find their origin in WZ production. Half of the remaining 20% is due to nonprompt events, mainly from DY . The results of this validation where predicted and observed events are compared can be found in Figure 7.22. As good agreement between the yields is observed, no further normalisation factor is derived. We do, however, apply an uncertainty of 10% to this background type based on the level of agreement between data and simulation in the kinematic shapes observed in this control region.

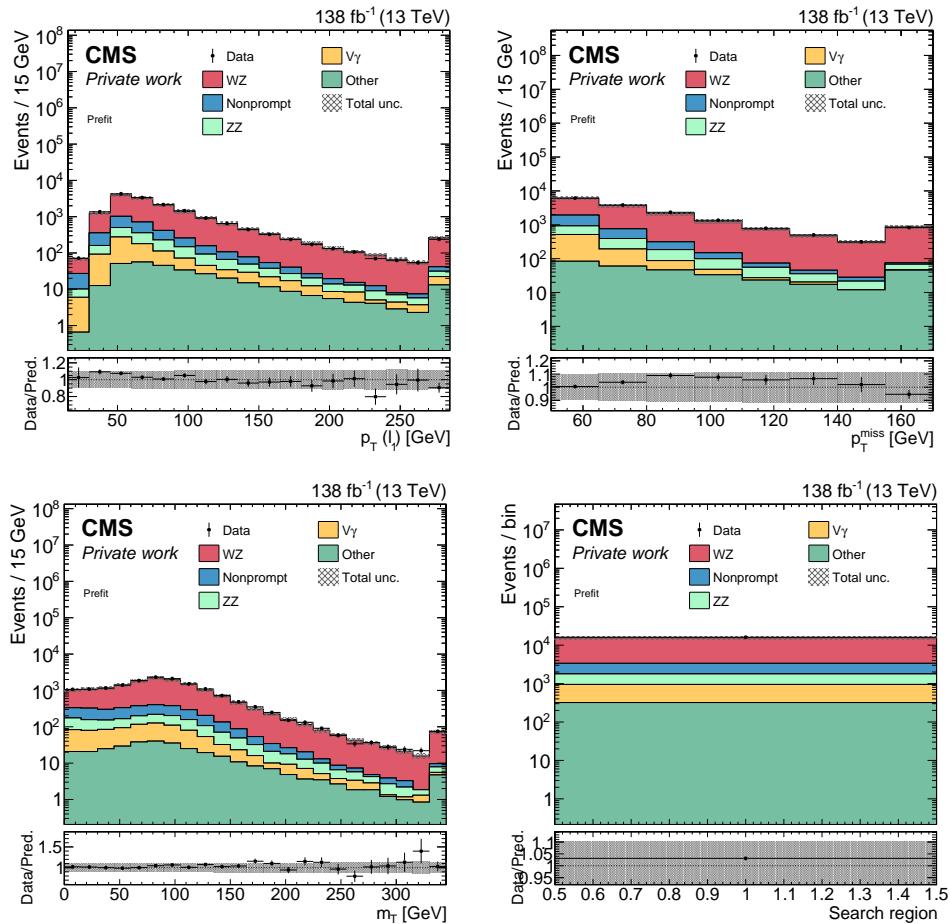


Figure 7.22: Comparison of the predicted (coloured histograms) and observed (points) yields in the WZ control region. Shown as a function of the leading lepton p_T (upper left), the missing transverse energy (upper right) and the transverse mass of the lepton not included in $\min m(2\ell|\text{OS})$ (lower left). The lower right plot shows the total yield.

7.8.3 ZZ background

Because of the presence of the fourth FO veto in the baseline selection, a fair amount of ZZ production is removed early on. The remaining collection of ZZ events in the background mainly comes from events where misreconstruction was present or where one of the leptons simply did not pass the selection. Their contribution to the total background yield is estimated from simulated samples. Similar to the WZ background, a validation is performed in a

dedicated control region which is defined as follows:

- exactly four tight light leptons with $p_T > 15, 10, 10, 10$ GeV;
- no fifth FO light lepton;
- the four light leptons form two OSSF pairs with $|m(2\ell|\text{OSSF}) - m_Z| < 15$ GeV;
- no loose b-jets;
- for each OS light lepton pair, $m(2\ell|\text{OS}) > 12$ GeV (to remove low-mass resonances).

If there are four leptons of the same flavour in the event, the two OSSF candidates Z_1 and Z_2 are chosen in such a way that the mass difference between each pair and the Z mass is minimised. In any case, Z_1 is defined as the OSSF pair with mass closest to the Z mass. The purity of ZZ in this control region is close to 97%. An underprediction was observed in this control region and therefore a scale factor of 1.12, derived as the ratio of the total yields in data and simulation, is applied to all simulated ZZ samples. The comparison of expected and observed yields in this region after application of this scale factor can be found in Figure 7.23. Good agreement is observed and an uncertainty of 10% is applied to cover any residual difference between data and simulation.

7.8.4 Background from photon conversions

Asymmetric photon conversion background, especially $Z\gamma$, is a non-negligible contribution to the analysis. Mainly in the entire low-mass region and the phase space of low $\min m(2\ell|\text{OS})$ and $M_{3\ell}$ in the high-mass region, it is an important contributing factor to the total background yield. It is once again estimated from simulation and validated in a $Z\gamma$ -enriched control region. This control region is based on the presence of three leptons with an invariant mass that is close to the Z boson mass. It is defined as follows:

- exactly three tight light leptons with $p_T > 15, 10, 10$ GeV;
- no fourth FO light lepton;
- trilepton invariant mass $|m(3\ell) - m_Z| < 15$ GeV (to select $Z\gamma$ events where the photon is radiated from one of the Z leptons);
- two leptons have OSSF and $|m(2\ell|\text{OSSF}) - m_Z| > 15$ GeV (to remove nonprompt DY and WZ contributions);
- no loose b-jets.

The purity of this selection is 70%. Half of the remaining events are due to nonprompt contributions, similarly to the WZ control region. It was found that the observed yield was larger than the simulated prediction. Similarly to the ZZ simulation, we apply a scale factor of 1.11 to all simulated $Z\gamma$ samples. The results of this control region after applying this scale factor can be found in Figure 7.24. Good agreement is observed and again a systematic uncertainty of 10% is applied to cover any residual differences between data and simulation.

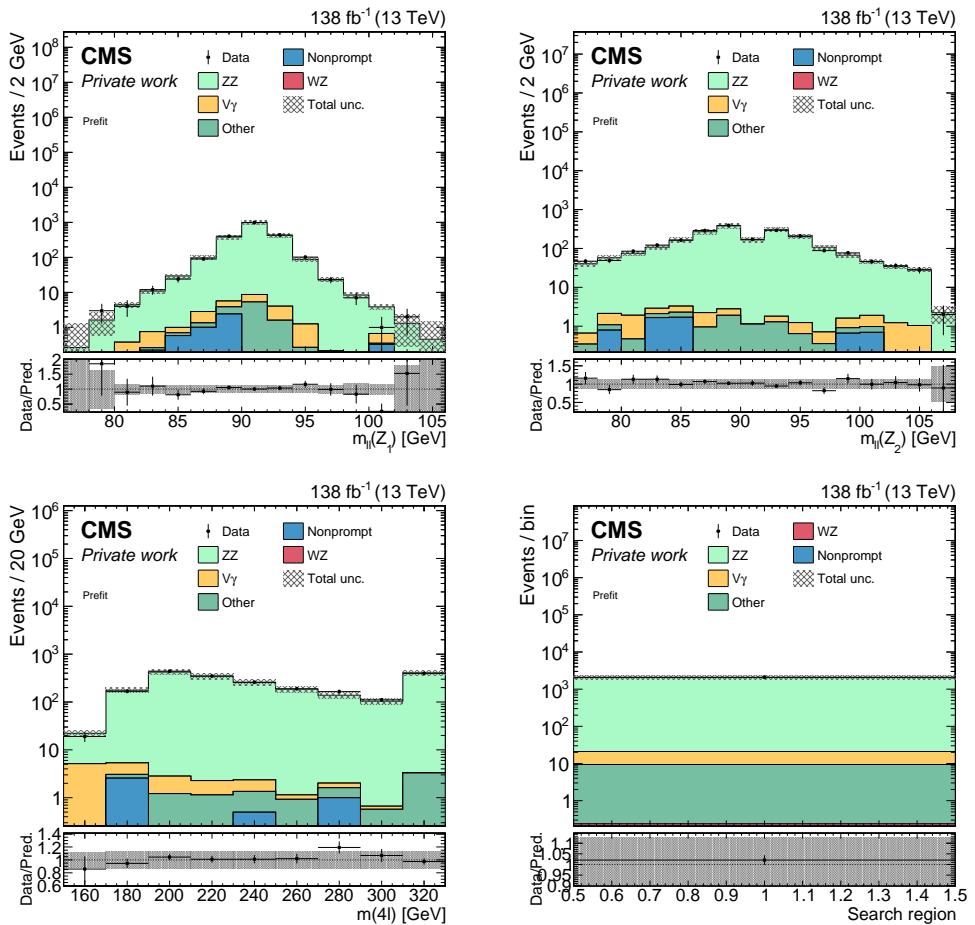


Figure 7.23: Comparison of the predicted (coloured histograms) and observed (points) yields in the ZZ control region. Shown as a function of the invariant mass of the lepton pair closest to the Z boson mass (upper left), the invariant mass of the other lepton pair (upper right) and the invariant mass of the 4 lepton system (lower left). The lower right plot shows the total yield.

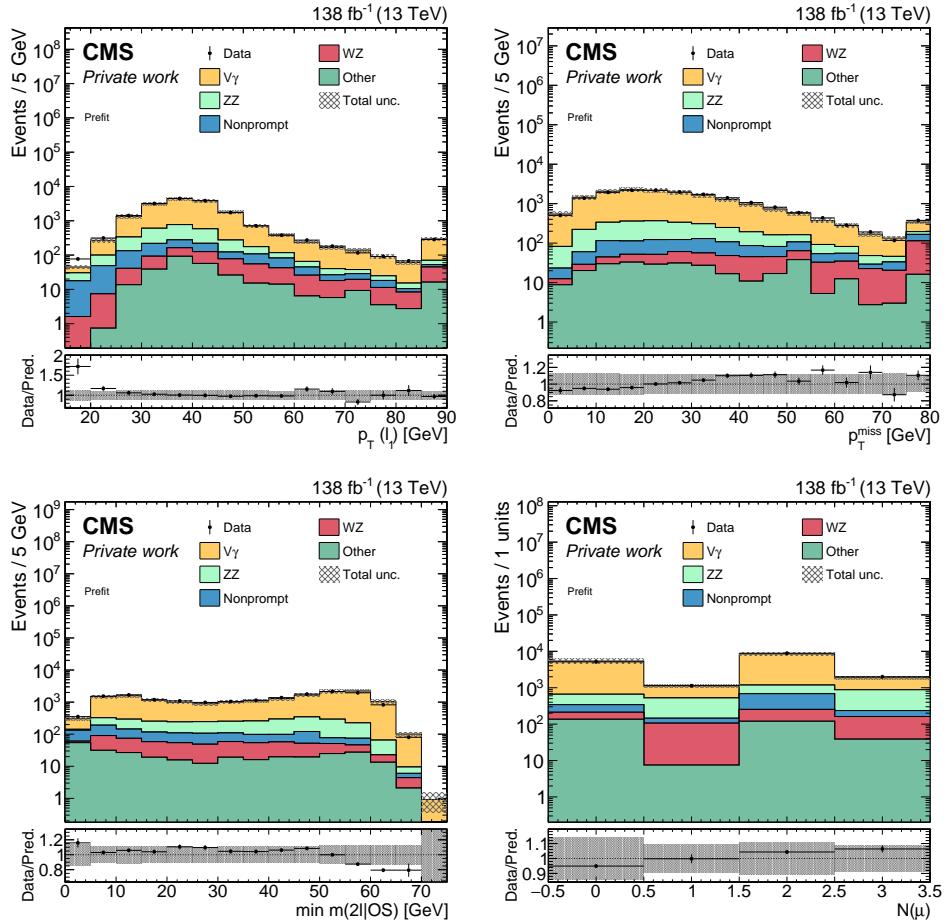


Figure 7.24: Comparison of the predicted (coloured histograms) and observed (points) yields in the conversion control region. Shown as a function of the leading lepton p_T (upper left), the missing transverse energy (upper right), $\min m(2\ell|OS)$ (lower left) and the number of muons in the event (lower right).

7.8.5 Charge misidentification background

In most of the search regions, background contributions that are due to the charge of an electron being misidentified are rare. The main exception is found in search region Ha (see Table 7.10) for final states with two same-sign electrons and a muon where additional selections are performed to filter out most of this background type. In order to estimate the surviving contributions, predictions from simulations where the charge of one of the reconstructed electrons is not the same as its generator match are used. However, past analyses have shown there to be discrepancies between simulation and data when it comes to charge misidentification in 2017 and 2018 [199, 203]. Therefore, we apply scale factors to all such events with the following value: 1.03 in 2016, 1.48 in 2017 and 1.34 in 2018. They were obtained by performing closure tests in a DY-enriched region using large data and simulation samples. As such, the statistical uncertainty on these numbers is negligible.

7.9 Systematic uncertainties

There are quite a few similarities between this analysis and the SUSY analysis, which causes the systematic uncertainties in both analyses to be comparable. The systematics on the luminosity, pileup and b-tagging use the exact same methodology as discussed in Section 6.8. Their exact values, however, might be different from the ones quoted in the last chapter due to updated calibration and validation in the UL samples. Because the other systematic uncertainties were estimated using methodologies that differ from the ones described in Section 6.8, we will briefly go over them below:

- The uncertainty on the **trigger efficiency** is once again split into two components: one to quantify the statistical limitations of the efficiency measurement and one to quantify the systematic component of the method. The first is uncorrelated between the different years of data-taking and is obtained from the statistical uncertainty on the ratio between the efficiency in data and simulation as evaluated at the p_T of the leading light lepton. The second is obtained from the difference between the aforementioned ratio and 1, which results in a value of 3%.
- Systematic uncertainties related to the **light lepton reconstruction, identification and isolation** corrections are applied in the analysis. For the electron reconstruction scale factors, the accompanying up and down variation weights are applied as a systematic uncertainty in the fit. As no reconstruction corrections are applied to muons, there are no reconstruction systematics for this lepton flavour either. The effect of the variations to the electron and muon energy scales was studied and was found to be negligible for the electron smearing and muon momentum scale systematics. The variations for the electron scaling were small but non-negligible. Therefore, no muon scale and electron resolution systematics are implemented and a flat 0.3% uncertainty per electron in the event is applied to account for the electron scaling. Lastly, the muon and electron identification and isolation uncertainties come from the measurement of the scale factors linked to the lepton MVA. They are determined and applied separately for muons and

electrons and done so as a function of light lepton p_T and $|\eta|$. While a single set of uncertainties take into account both identification and isolation, a split is made between the statistical limitation of the scale factor measurement and the systematic uncertainty of the method. The latter is the dominant contributor to the uncertainty.

- The DeepTau algorithm also comes with systematic uncertainties on τ_h **identification and isolation** that were obtained using the tag-and-probe method. They come in the form of scale factors for up and down variations to be applied per genuine τ_h , electron faking τ_h signatures and muon faking τ_h signatures in the event. The uncertainty on the τ_h energy is applied by varying the $\tau_h p_T$ of all candidates in simulation up and down and propagating the changes through the analysis. Usually, this is done for both genuine τ_h and τ_h faked by electrons. As the impact of the latter turned out to be negligible, there is no τ_h scale uncertainty applied for this type of τ_h candidate. The energy scale uncertainty for genuine τ_h is applied, however, and is taken to be uncorrelated across years and decay modes.
- The analysis is sensitive to leptons but not as much to jets. Therefore, we do not apply the full **JEC, JER and unclustered MET** uncertainties that warrant varying all relevant variables up and down one by one and rerunning the analysis. This strategy was studied but yields incoherent fluctuations in the variations. Instead, two flat 3% uncertainties are applied, one of which is uncorrelated across all years of data-taking while the other is correlated.
- For the **nonprompt lepton background** systematics, a slightly different approach is used between the three flavours, which are all decorrelated from each other. The data-driven nonprompt τ_h background contribution receives a 30% flat uncertainty, which is derived from the goodness of closure in the closure tests. The same flat uncertainty is applied to muons as well. In case of electrons, the situation is slightly different, however. Because there is a clear structure in electron closure as function of leading light lepton p_T , it is instead opted to make the systematic uncertainty on nonprompt electrons dependent on this variable. The exact values of the systematic uncertainty applied are 15% for $p_T < 35$ GeV, 30% for $35 < p_T \leq 55$ GeV and 50% for $p_T \geq 55$ GeV. Because the nonprompt τ_h sources in events with and without an OSSF light lepton pair tend to be different, the systematic uncertainties on nonprompt τ contributions are decorrelated between those two types of events in order to better constrain the nonprompt contributions.
- All **prompt backgrounds** are predicted using simulation. In the case of background from ZZ processes, WZ processes and conversions, these predictions were validated in dedicated control regions. There, data and simulation were found to agree within 10%. Therefore, a flat 10% uncertainty is applied to these three background sources. The other prompt backgrounds do not have a dedicated validation region. Instead, their uncertainties are obtained from literature. In the case of backgrounds with three SM bosons these come from SM predictions of their production [204–207], which have

a precision of 10%. Measurements of $t\bar{t}X$ processes in CMS show that the ones of interest to this analysis have an experimental precision of 10% [203, 208, 209], which is consequently the value assigned to the systematic uncertainty of this background category. All the remaining SM processes that contribute little to the analysis are bundled together. Vector-boson-associated H boson production is the main contributor to this collection and therefore the systematic uncertainty is set to the corresponding 20% experimental uncertainty [5].

- The main source of systematic in the **charge misidentification** background comes from the scale factors that are applied to the simulation. It is addressed by applying a flat 15% uncertainty on the relevant background contribution. This number is obtained from closure tests performed in previous analyses [199, 203].
- It is not just the background prediction that comes with uncertainties. The **theory prediction** of the HNL processes comes with uncertainties on the theoretical calculations and event simulation as well. There are a few main contributing factors that need to be considered. The first is due to the renormalisation and factorisation scales (μ_F and μ_R) that are needed in the QCD calculations to avoid singularities. They are estimated by varying their values up and down by a factor two independently and taking the envelope of maximum variations [210]. Secondly, the choice of the PDF will come with uncertainties. To estimate them all parameters in the PDF are varied and the root mean squared (RMS) is calculated [173]. For the signal models with $m_N < m_W$, we used LO samples, which means they come with large uncertainties of up to 20% on the factorisation and renormalisation scales. In order to reduce these uncertainties, LO-to-NNLO K-factors for $W + \text{jets}$ production are evaluated and applied [96]. Their value lies at 1.089 and allows the uncertainty to be contained to a flat 4% instead. In the mass region with $m_N > m_W$, no such K-factors are applied as the samples are generated at NLO. A study was performed instead to look at the effect of the renormalisation and factorisation scales and the PDF choice. The results can be found in Figure 7.25, which shows that for DY production of HNLs, the scale variations are the dominating systematic while VBF production is much more sensitive to the PDF choice. The conclusion was then to apply a 3% flat uncertainty due to scale variations for the first type of production and to apply a 15% flat uncertainty due to PDF choice to the latter.
- Over the course of data-taking in Run II, and especially in 2017, an issue with the L1 trigger caused it to throw away events unjustly every now and again. This issue is called the **prefiring issue** and was found in both the electron/photon and muon sector. Over time the crystals used in the ECAL systems get darker. This causes the trigger tower to be assigned to the previous bunch crossing by accident. As the probability of this happening depends on the opacity of the crystals, this issue worsened over time. In the muon system, similar effects are caused by the limited time resolution. To account for it, correction weights are applied to all simulations to emulate the issue. The uncertainty that comes with these weights is used in the analysis as a systematic uncertainty.

An overview of all the systematics used in the final fit is given in Table 7.14.

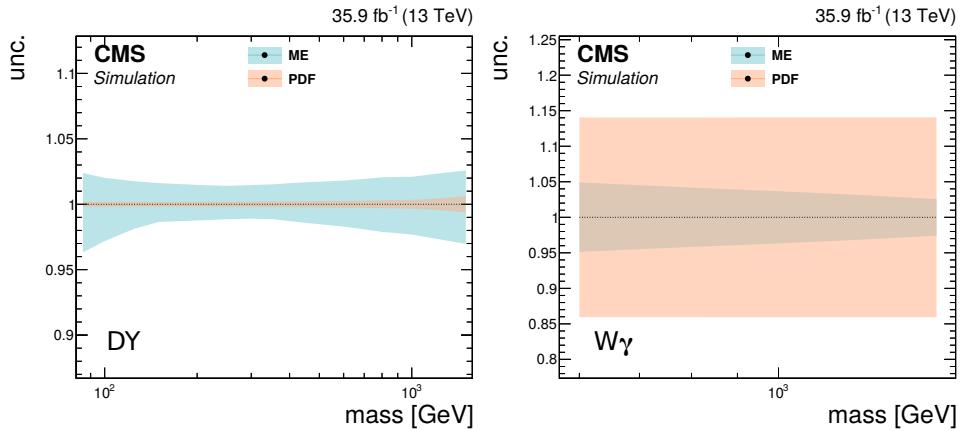


Figure 7.25: Comparison of the renormalisation-factorisation scale systematics (ME) and the PDF systematics as measured in an HNL sample with exclusive electron neutrino coupling with DY production (left) and VBF production (right). The envelope corresponding to ME systematics is obtained by taking the maximum of all fluctuations while the one corresponding to the PDF systematic is obtained from the RMS of all variations.

The impacts of the systematic uncertainties with the highest impact in the analysis are shown in figure 7.26 for a fit to signal plus background in the low-mass ($m_N = 60$ GeV, exclusive electron neutrino coupling) and high-mass ($m_N = 200$ GeV, exclusive τ neutrino coupling) signal region. The impact Δr is defined here as the change in signal strength when a nuisance parameter is varied up and down by one standard deviation. The plots also show the pull of each nuisance parameter in the fit, referring to the fractional difference between the value of the nuisance parameter after the fit and its nominal value. The error on the pull is a measure for how much a nuisance parameter is constrained in the fit. If it is much smaller than a range of 1, it means that the uncertainty was overestimated. From these figures it can be seen that for the low-mass region fit, the statistical uncertainties, described in the plot by entries that have a name beginning with “prop_bin”, have the largest impact. Other important systematics in the low-mass region consist of the nonprompt background normalisation, prompt background normalisation and pileup reweighting. In the high-mass region, one can see that while statistical uncertainties have the largest impact, the importance of systematic uncertainties is larger compared to the low-mass region. The systematic uncertainty sources with the highest impacts are similar to those observed in the low-mass region. No irregular values are observed for the impacts, pulls and constraints in this analysis.

7.10 Results

With all the pieces in place, it is time to unveil the signal regions. The observed yields are plotted against the total predicted background yield and this is overlayed with the expected yield of a few relevant HNL signal models for all distributions relevant to the fit as discussed in Section 7.7.4. First in line is the **low-mass region**. The yields as a function of the signal regions and BDT output scores relevant to the fit strategy of that region as defined in

Table 7.14: Systematic uncertainties considered in the analysis, their typical size and the processes they are applied on.

| Source | Typical size | Processes applied on |
|-----------------------------|--------------|------------------------------------|
| Luminosity | 1.6% | Background and signal MC |
| Trigger efficiency | 2% | Background and signal MC |
| WZ normalisation | 10% | WZ background |
| ZZ normalisation | 10% | ZZ background |
| Conversion normalisation | 10% | Conversion background |
| Charge misid. normalisation | 15% | Charge misid background |
| t̄X normalisation | 10% | t̄X background |
| Triboson normalisation | 10% | Triboson background |
| Other normalisation | 20% | “Other” background |
| Nonprompt τ_h | 30% | Nonprompt background |
| Nonprompt light lepton | 30% | Nonprompt background |
| Pileup | 2–4% | Background and signal MC |
| Electron reconstruction | 2–3% | Background and signal MC |
| Electron identification | 1–2% | Background and signal MC |
| Muon identification | 1–2% | Background and signal MC |
| Tau identification | 1% | Background and signal MC |
| Tau energy scale | <1% | Background and signal MC |
| b-tagging | 1–3% | Background and signal MC |
| JEC, JER, unclustered MET | 3% | Background and signal MC |
| Prefiring | <1% | Background and signal MC |
| Theory: K -factor | 4% | signal MC: DY prod. $m_N < 80$ GeV |
| Theory: μ_F, μ_R | 3% | signal MC: DY prod. $m_N > 80$ GeV |
| Theory: PDF choice | 15% | signal MC: VBF prod. |

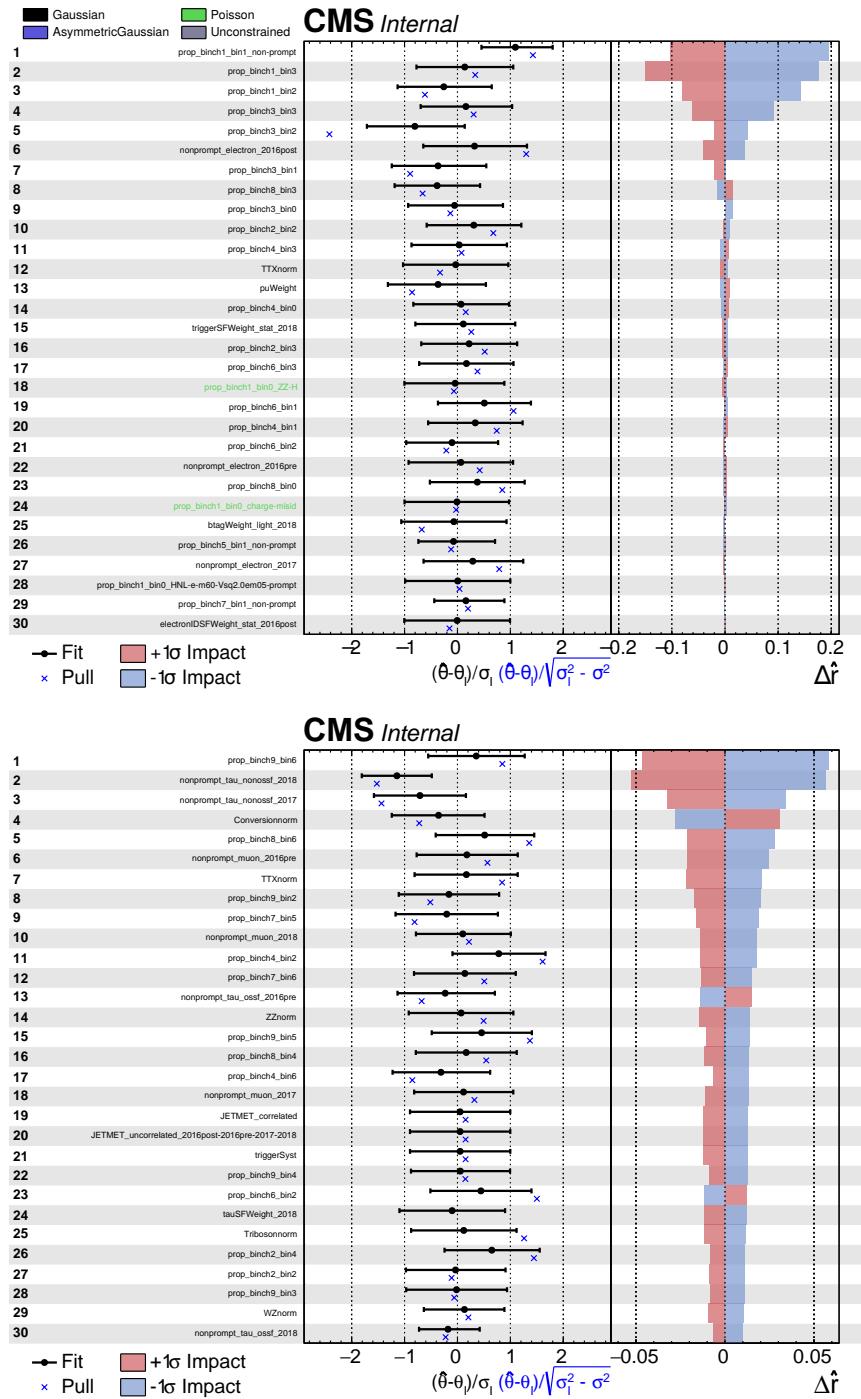


Figure 7.26: Impacts and pulls for a fit to a signal with $m_N = 60$ GeV and exclusive coupling to an electron neutrino plus background in the low-mass region (top) and a fit to a signal with $m_N = 200$ GeV and exclusive coupling to a τ neutrino plus background in the high-mass region (bottom).

Table 7.12 are shown in Figures 7.27 to 7.38 for both the fully light lepton final states and the final states containing one τ_h for all three coupling scenarios. In these plots, the predicted background yields are shown with their best fit normalisation from a background-only fit. Good agreement between data and SM background prediction is observed and none of the bins show a significant excess.

Secondly, we take a look at the **high-mass region**. The resulting yields for the distributions relevant to the fit as defined in Table 7.12 are depicted in Figures 7.32 to 7.37 for all three coupling scenarios. In general, good agreement is found between data and background prediction. The only exception can be found in the output score of $\text{BDT}(300\text{--}400, e, 0\tau_h)$ in events with a final state of two same-sign electrons and a muon, where the formation of a slight excess can be observed as can be seen in Figure 7.32. This excess has been investigated and was found to not be significant with a local significance of 2.88σ and a global significance of 2.25σ . The local significance was measured using only the data in the relevant $\text{BDT}(300\text{--}400, e, 0\tau_h)$ distribution while the global significance takes the data from all search regions into account.

7.10.1 Interpretation

With no significant excess in the observed yields, we find ourselves in the same situation as the previous HNL analysis performed at UGent with the 2016 dataset. Exclusion limits can now be calculated by performing simultaneous fits of the relevant distributions using the Asymptotic approximation of the CL_s method with the LHC test statistic [139, 211].

This fit results in an upper limit value of the signal strength of the signal process in question. If this signal strength value is found to be smaller than 1, it means the signal process at that cross section is excluded at 95% CL. Scaling the input cross section with the signal strength allows us to find the upper exclusion limit on the cross section of that HNL process. As demonstrated in Figure 3.5, exclusion limits on HNL models are typically given in the $|V_{\ell N}|^2\text{-}m_N$ parameter space. Given the proportional nature between $|V_{\ell N}|^2$ and the cross section and the fact that in prompt scenarios there is nearly no dependence of the acceptance on the exact value of the coupling, we can simply translate the interpretation of the signal strength to the coupling squared. The signal strength can be extrapolated along an inverse function of $|V_{\ell N}|^2$ and thus the exclusion limit for a certain mass point can then be found by scaling the input $|V_{\ell N}|^2$ by the signal strength. In the case where the HNL is displaced, the situation is a bit less straightforward. As changing the input coupling squared will change the mean lifetime of the HNL, the acceptance of prompt events will be influenced as well. The relation between signal strength and coupling squared will no longer be a simple inverse function and instead of fitting a single coupling-mass point we will have to perform fits for a range of coupling values at a fixed mass point until we find the coupling squared value where the signal strength is one. In practice, this is done by starting from a sample with fixed mass and fixed coupling squared and using the reweighting function in Section 7.3.1 to obtain input distributions at different coupling values.

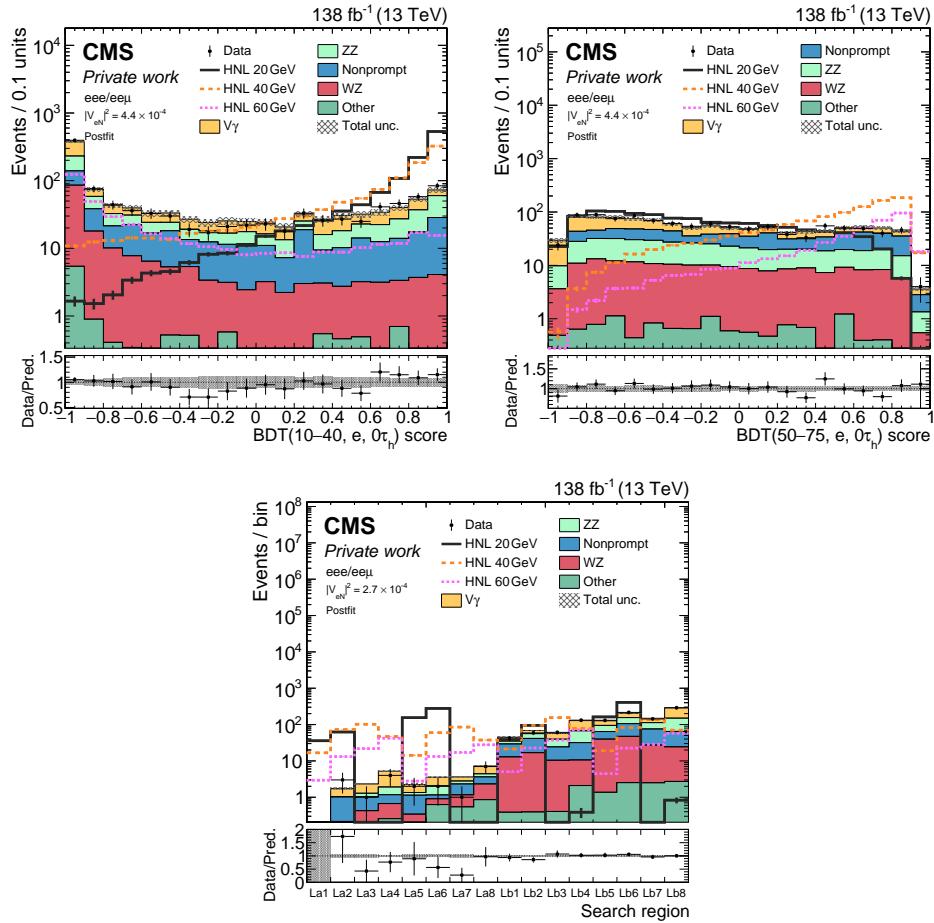


Figure 7.27: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to v_e (coloured lines) corresponding to three different mass assumptions in the low-mass region selection with final states containing at least two electrons using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show the output score of BDT(10–40, e, $0\tau_h$) (top left), the output score of BDT(50–75, e, $0\tau_h$) (top right) and search regions La1–8 (bottom). All yields displayed in the output score distributions need to be contained in search regions Lb1–8.

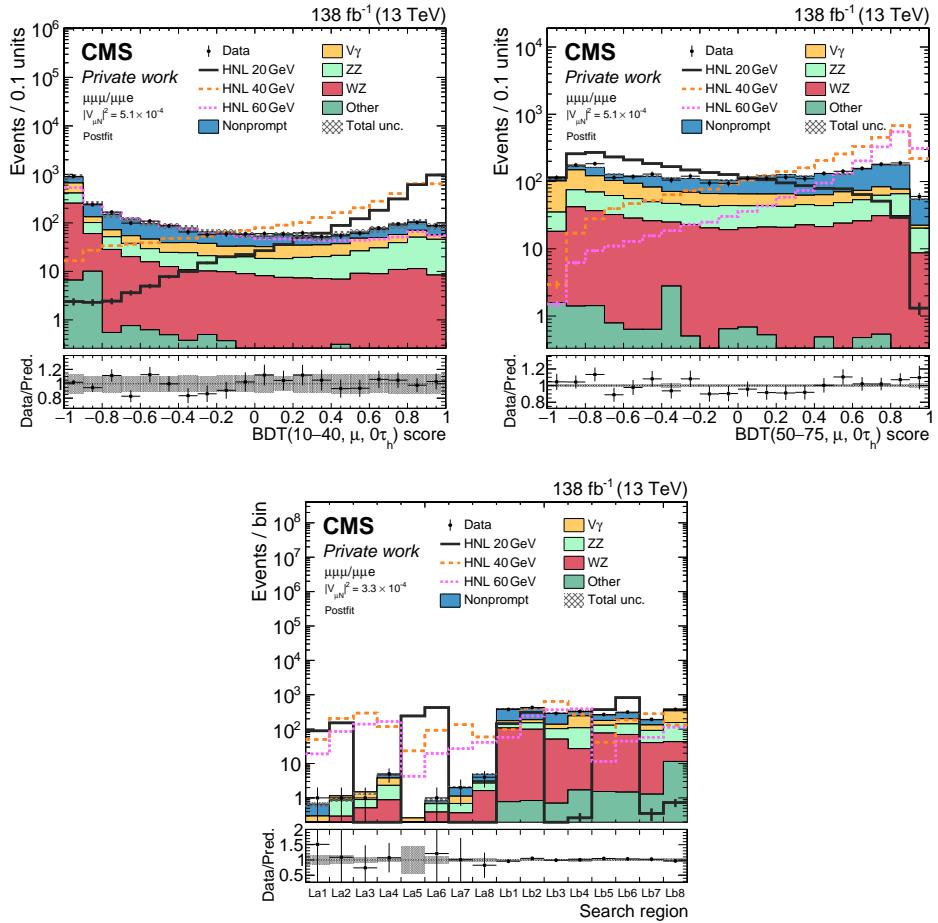


Figure 7.28: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_μ (coloured lines) corresponding to three different mass assumptions in the low-mass region selection with final states containing at least two muons using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show the output score of $BDT(10-40, \mu, 0\tau_h)$ (top left), the output score of $BDT(50-75, \mu, 0\tau_h)$ (top right) and search regions La1–8 (bottom). All yields displayed in the output score distributions need to be contained in search regions Lb1–8.

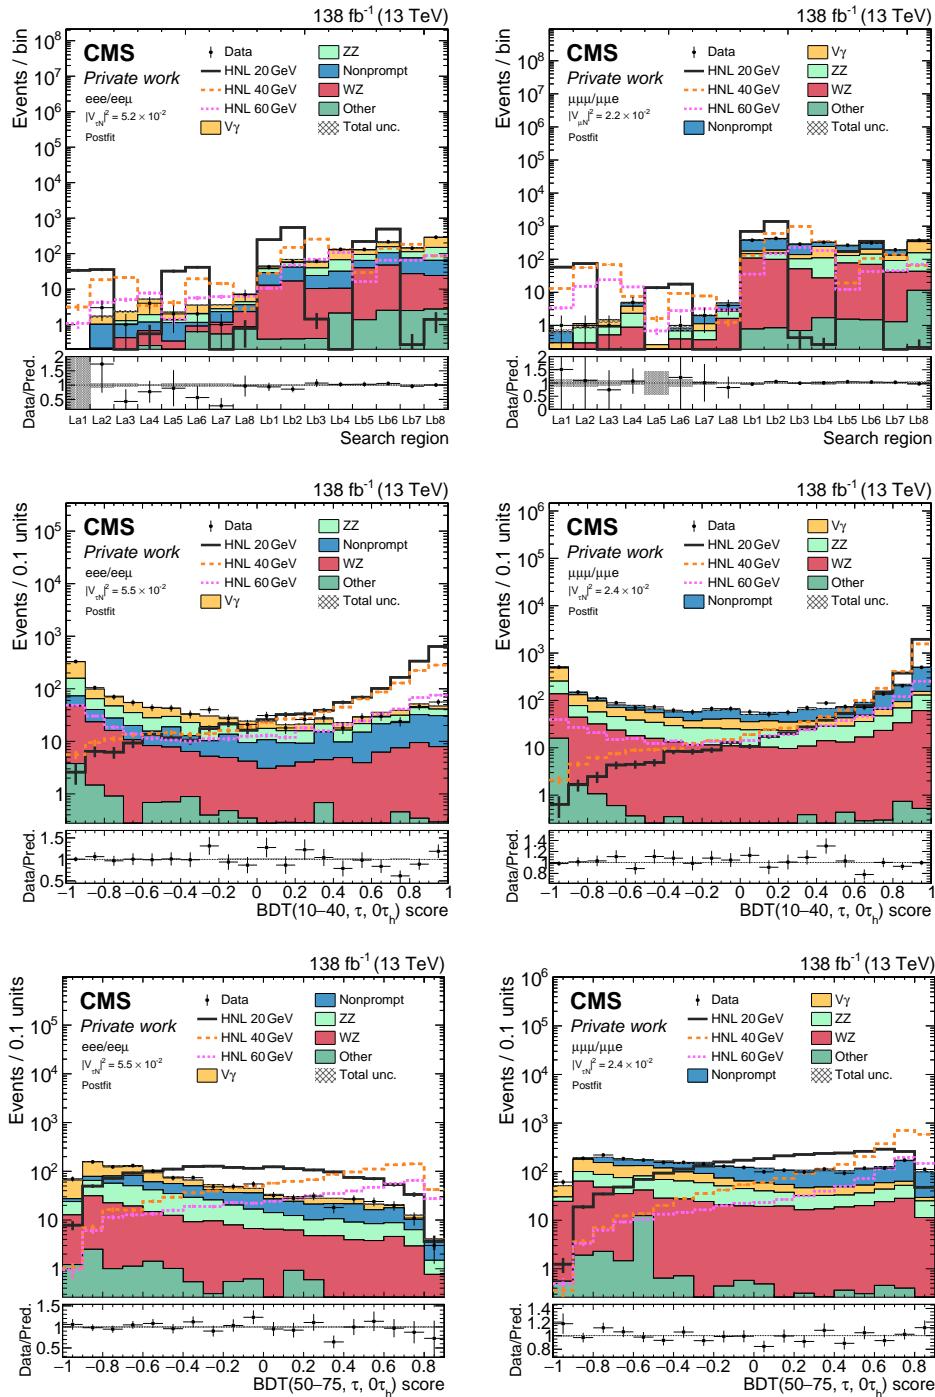


Figure 7.29: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_τ (coloured lines) corresponding to three different mass assumptions in the low-mass region selection with final states containing at least two electrons (left) or at least two muons (right) using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show search regions La1–8 (top), the output score of BDT(10–40, τ , $0\tau_h$) (middle) and the output score of BDT(50–75, μ , $0\tau_h$) (bottom). All yields displayed in the output score distributions need to be contained in search regions Lb1–8.

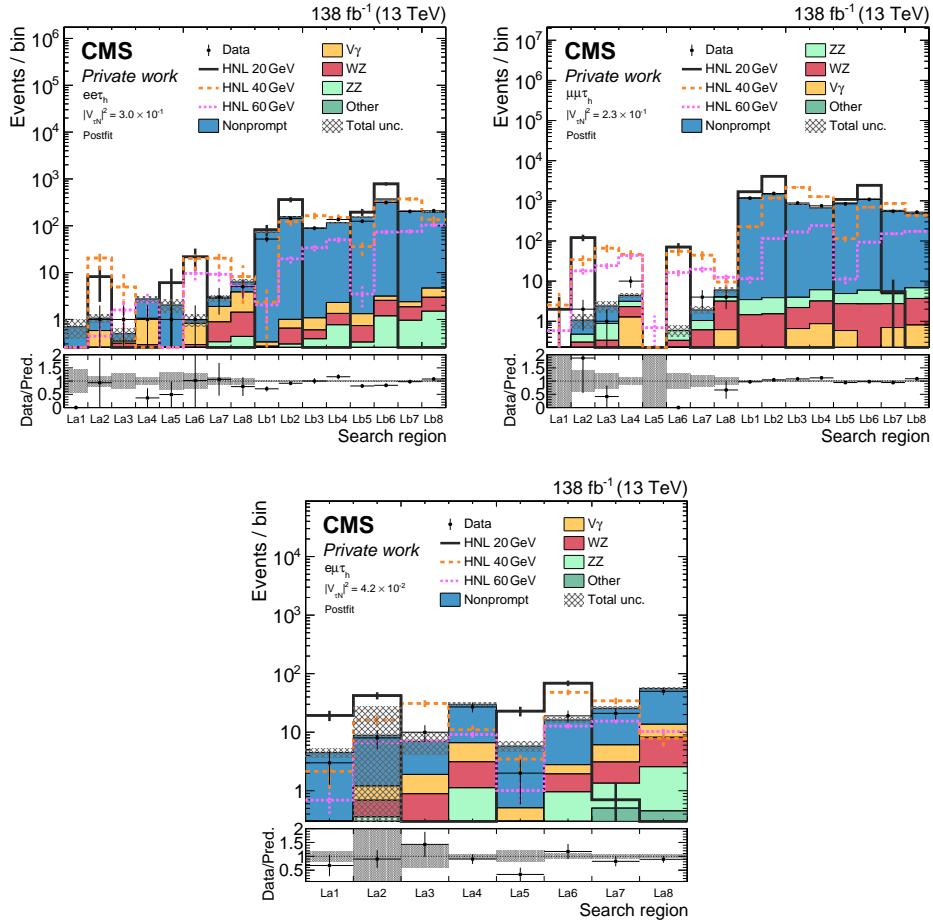


Figure 7.30: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_τ (coloured lines) corresponding to three different mass assumptions in the low-mass region selection with final states containing two electrons and a tau (top left), two muons and a tau (top right) or a muon, an electron and a tau (bottom) using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show search regions La1–8 for all three final states and Lb1–8 for the final states with two same flavour light leptons.

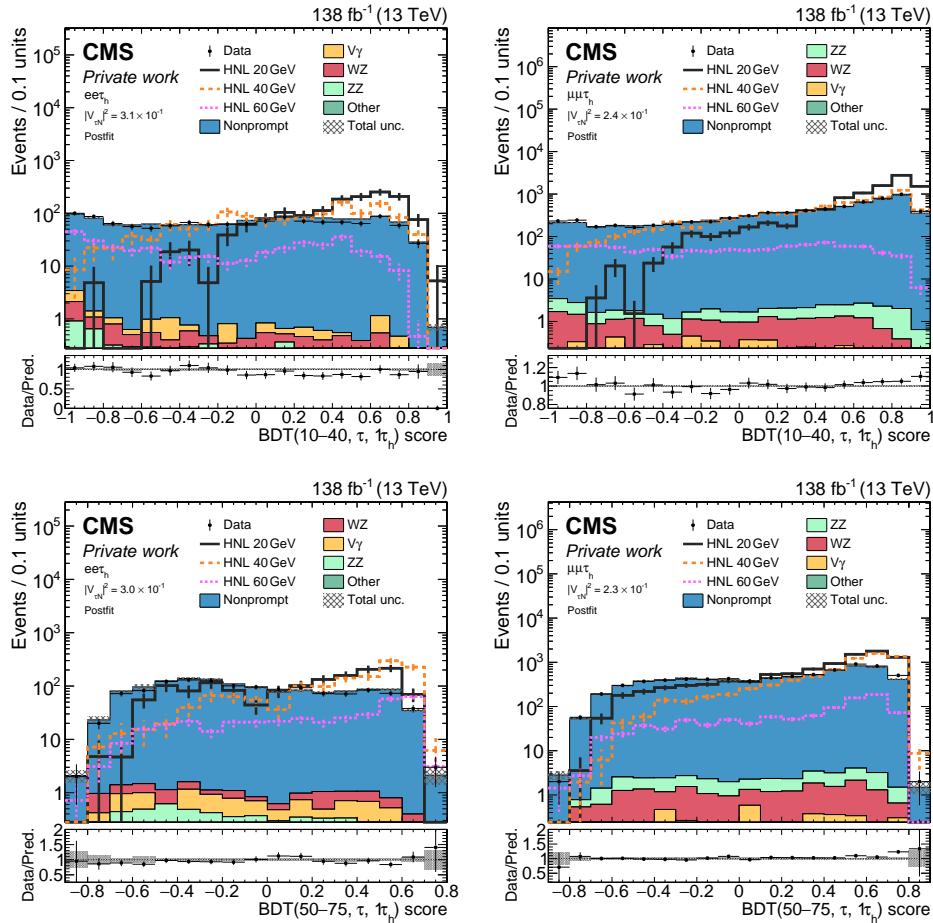


Figure 7.31: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_τ (coloured lines) corresponding to three different mass assumptions in the low-mass region selection with final states containing two electrons and a tau (left), and two muons and a tau (right) using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show the output score of $BDT(10-40, \tau, 1\tau_h)$ (top) and the output score of $BDT(50-75, \mu, 1\tau_h)$ (bottom). All yields displayed in the output score distributions need to be contained in search regions Lb1–8.

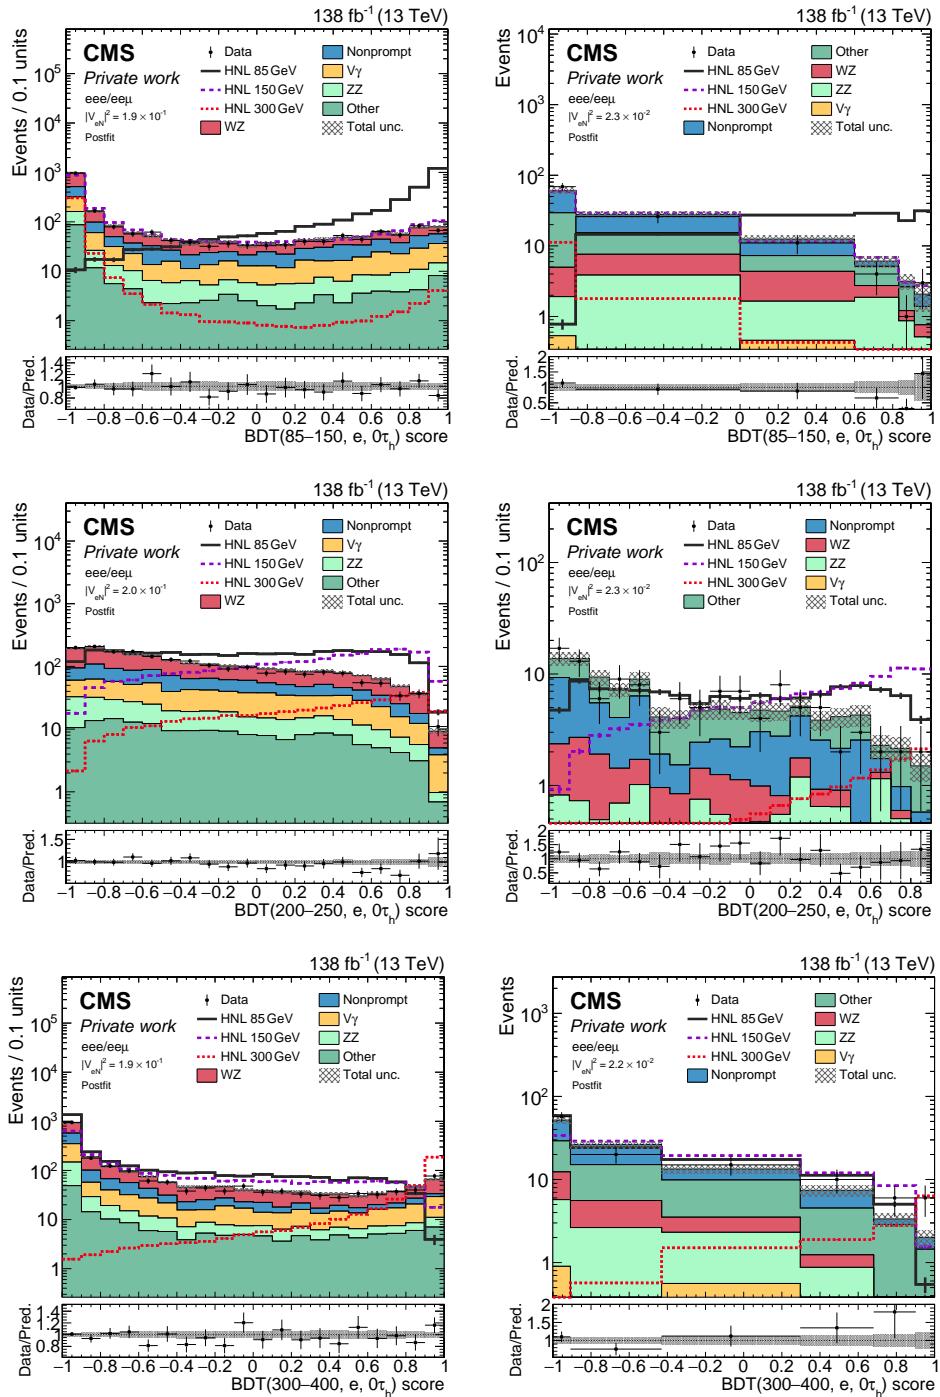


Figure 7.32: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to v_e (coloured lines) corresponding to three different mass assumptions in the high-mass region selection with final states containing at least two electrons using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show the output score of $BDT(85-150, e, 0\tau_h)$ (top), the output score of $BDT(200-250, e, 0\tau_h)$ (middle) and the output score of $BDT(300-400, e, 0\tau_h)$ (bottom). The left column contains only events with an OSSF pair while the right column contains only events without such a pair.

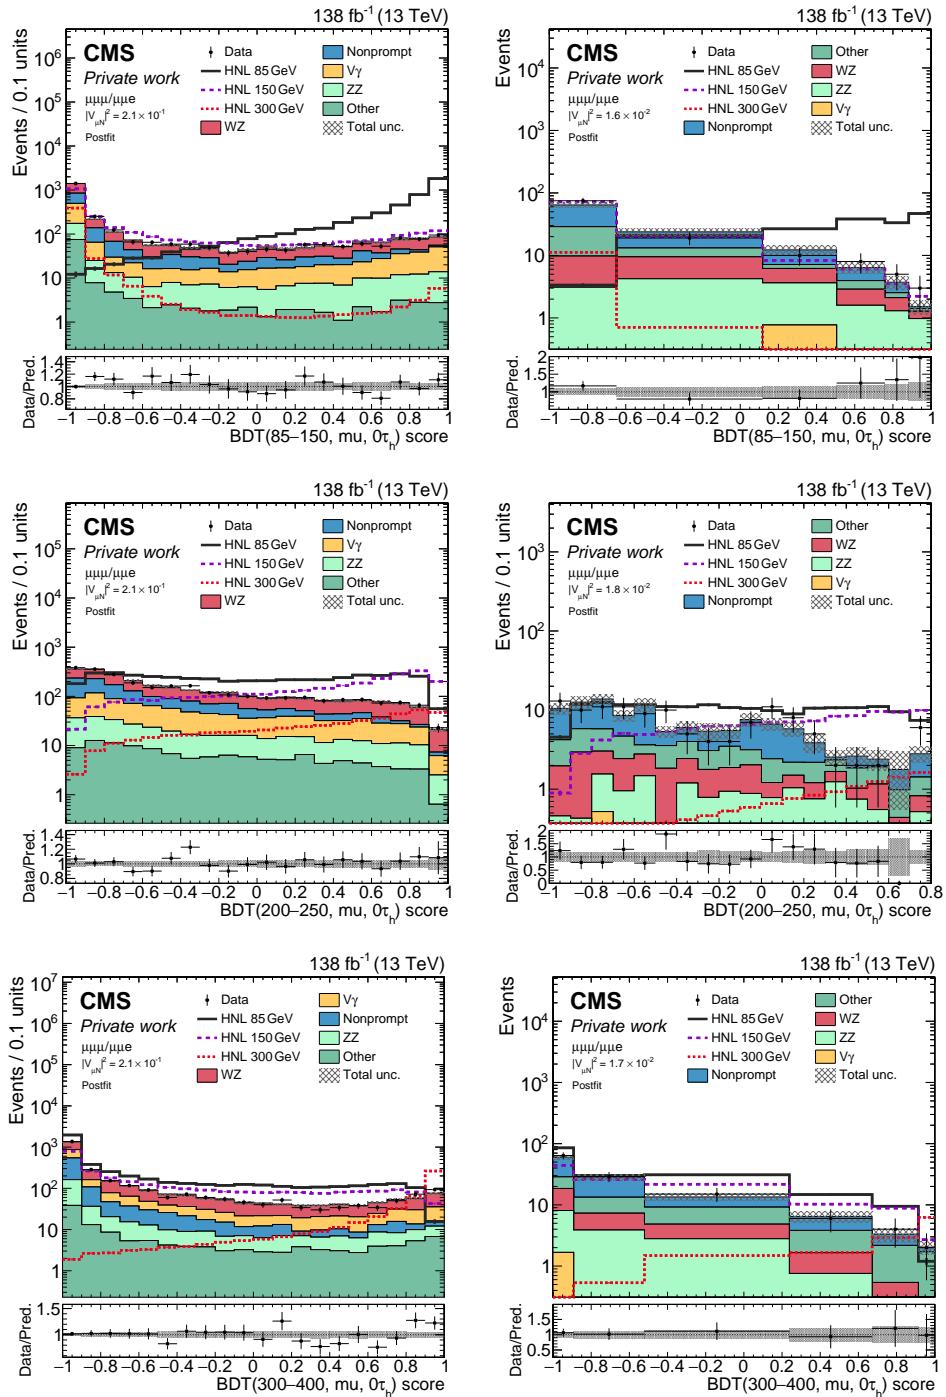


Figure 7.33: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_μ (coloured lines) corresponding to three different mass assumptions in the high-mass region selection with final states containing at least two muons using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show the output score of $BDT(85-150, \mu, 0\tau_h)$ (top), the output score of $BDT(200-250, \mu, 0\tau_h)$ (middle) and the output score of $BDT(300-400, \mu, 0\tau_h)$ (bottom). The left column contains only events with an OSSF pair while the right column contains only events without such a pair.

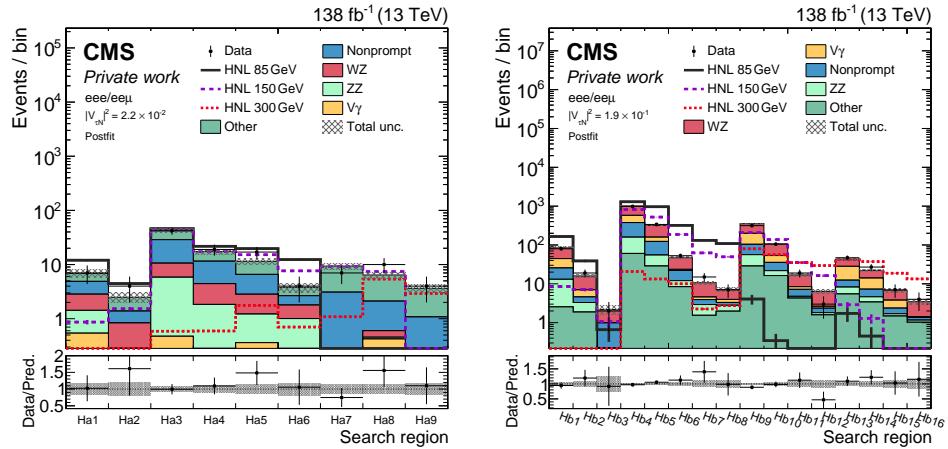


Figure 7.34: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_e (coloured lines) corresponding to three different mass assumptions in the high-mass region selection with final states containing at least two electrons using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show search region Ha (left) and search region Hb (right).

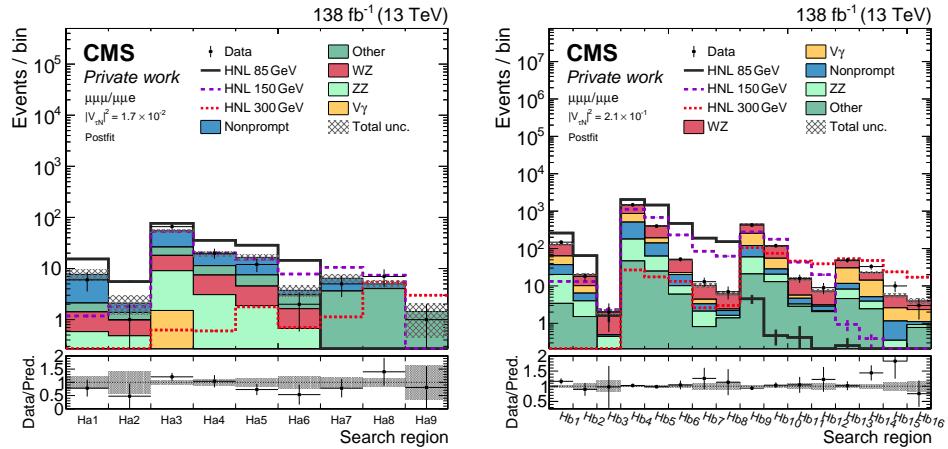


Figure 7.35: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_μ (coloured lines) corresponding to three different mass assumptions in the high-mass region selection with final states containing at least two muons using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show search region Ha (left) and search region Hb (right).

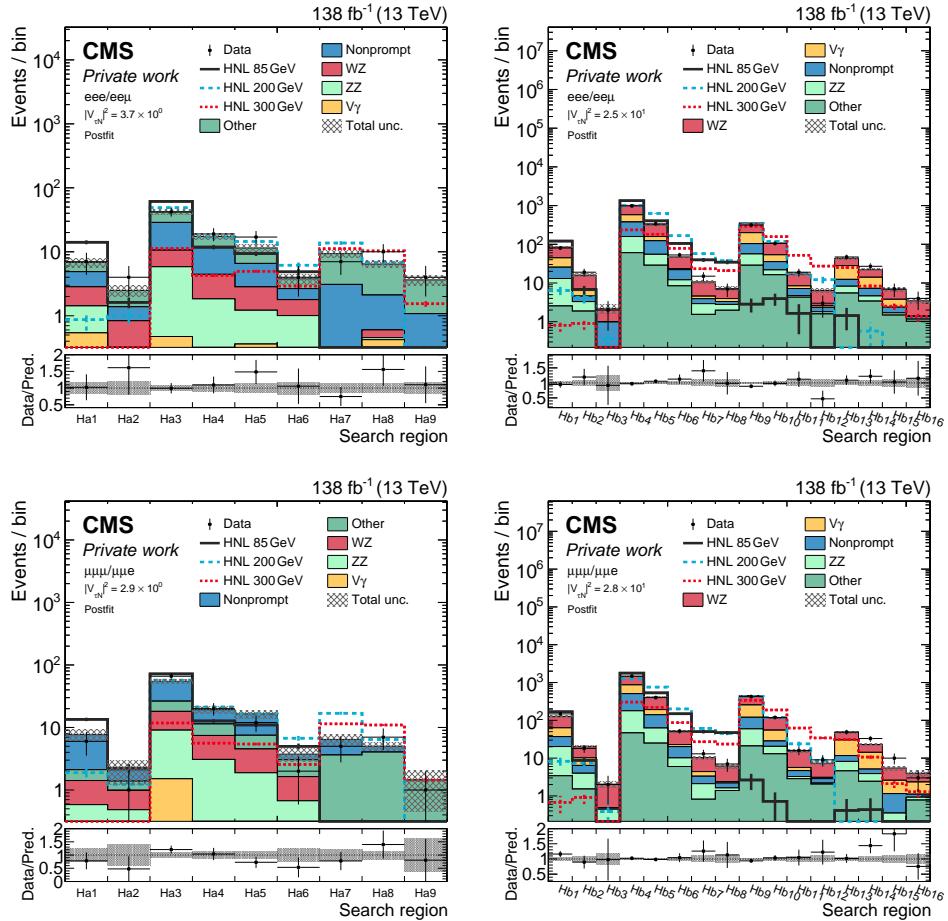


Figure 7.36: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_τ (coloured lines) corresponding to three different mass assumptions in the high-mass region selection with final states containing at least two electrons (top) or at least two muons (bottom) using Run II data. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show search region Ha (left) and search region Hb (right).

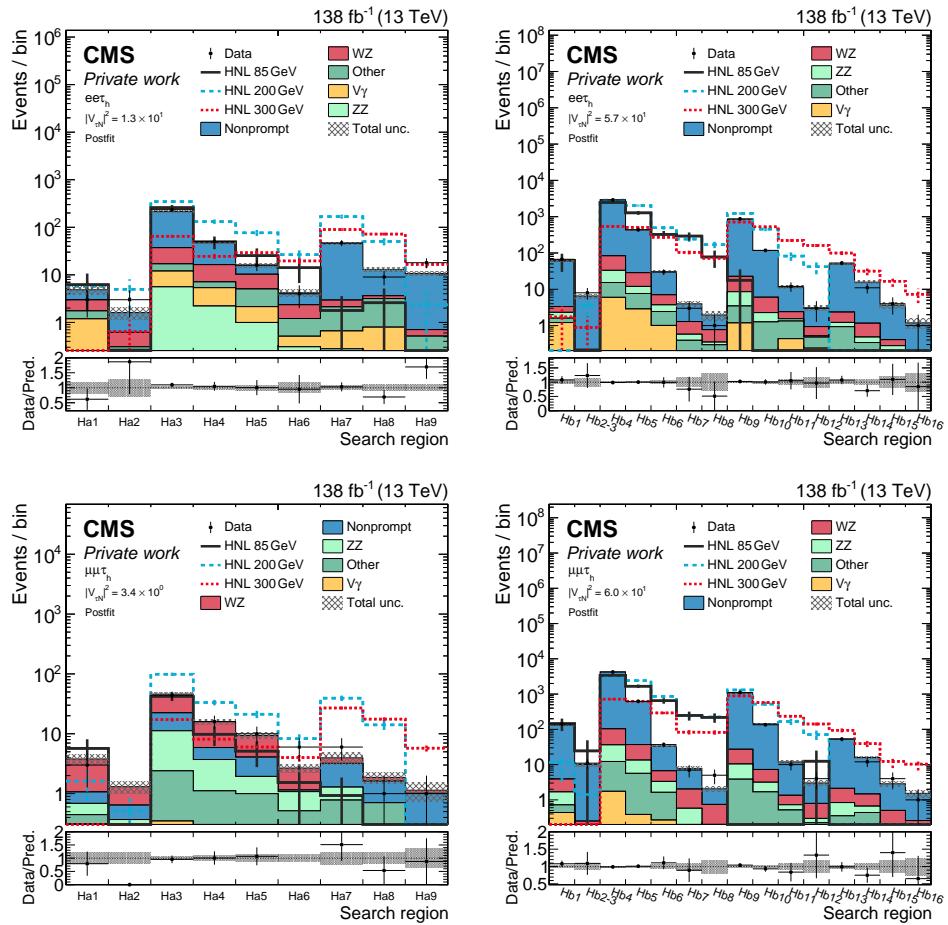


Figure 7.37: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_τ (coloured lines) corresponding to three different mass assumptions in the high-mass region selection with final states containing two electrons and a tau (top), and two muons and a tau (bottom). The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show search region Ha (left) and search region Hb (right).

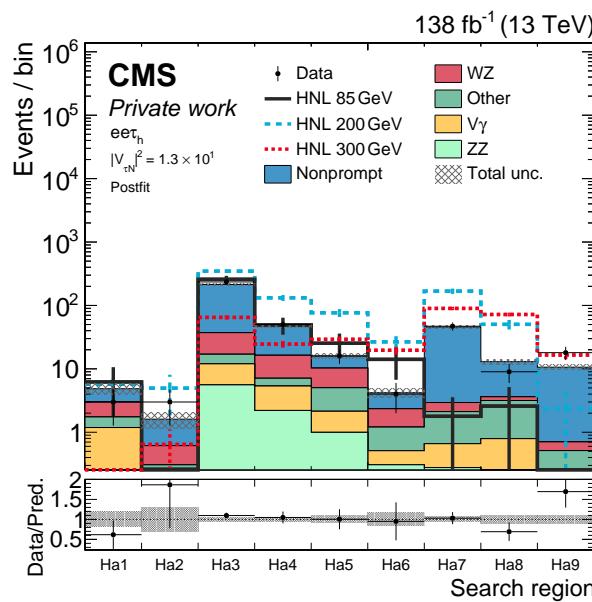


Figure 7.38: Expected background yields (coloured histograms), observed yields (points) and expected yields of HNL models with coupling to ν_τ (coloured lines) corresponding to three different mass assumptions in the high-mass region selection with final states containing a muon, an electron and a tau. The predicted background yields are shown with their best fit normalisation from the background-only fit. Vertical bars represent statistical uncertainties in data while hatched regions represent the combination of systematic uncertainty and statistical uncertainty in the predictions. The distributions show search region Ha.

A separate fit is performed for each mass point so that we have a collection of pairs of mass points and their respective exclusion limit on the coupling squared for the range of masses we consider. Such collections are created for each flavour scenario separately to obtain the limit plots presented later on. The exact mass points used are 10, 12.5, 15, 17.5, 20, 30, 40, 50, 60, 70, 75, 85, 100, 125, 100, 200, 250, 300, 350, 400, 450, 500, 600, 700, 800, 900, and 1000 GeV. For light lepton flavour coupling scenarios, additional mass points of 1200 and 1500 GeV are considered as well. Prompt samples are used for all mass points above 20 GeV while displaced samples are used for the other mass points. The input distributions to the fit are the ones as described in Section 7.7.4. On the border mass points between different search strategies, e.g. the 40 GeV scenario where we move to a different BDT score, a fit is performed for both strategies and both of these limits are quoted in the final result.

7.10.2 Exclusive light lepton neutrino coupling

In order to gauge what fraction of the expected improvement to the limits as compared to the 2016-only limits comes from the updated analysis strategy and what fraction comes from the increase in dataset size, expected limits using only the 2016 dataset with the updated strategy are studied. These expected limits assume that the observed data is equal to the predicted background. The result can be found in Figure 7.39. The dotted line shows the expected limit. The green band shows the 1 sigma uncertainty band on the expected limit while the yellow band shows the 2 sigma uncertainty band on the expected limit. Starting with the electron neutrino coupling scenario on the left-hand side, we can see that at low HNL mass, no big improvement is found with the updated strategy. The story is different for the scenario with exclusive muon neutrino coupling however with an improvement of a factor 2-3 at these low HNL masses. While it is true that part of the changes in the strategy is shared between all coupling scenarios, such as the BDT strategy, which would reflect in similar improvements, there were also major improvements in lepton identification which do reflect differently on these two scenarios. It can be concluded that the lepton identification improvements are a major player in the sensitivity boost, which makes sense if one remembers that in Ref. [99] the major background in the low-mass region originated from nonprompt leptons. On the other hand, the changes in the analysis strategy are mostly implemented in final states with a lot of background so it is not unexpected they will not have the largest impact on the final sensitivity.

In the high-mass region, we find similar improvements to the sensitivity in both the exclusive electron neutrino and the exclusive muon neutrino coupling scenarios. Especially in the medium mass range between 85 GeV and 400 GeV, the largest improvements are made. The origin of this increase in sensitivity can be traced back to the use of a completely new BDT strategy for these mass points. At higher masses, the cut-based search region strategy is still the default so here the improvement mainly comes from the improved lepton identification.

The observed 95% CL limits from Run II fits are shown as the full black line in Figure 7.40 in the scenarios of exclusive muon neutrino or electron neutrino coupling to a prompt Majorana HNL. This is compared to the results from the predecessor analysis that used only 2016

data [99]. At low masses and medium masses up to 300 GeV, we observe improvements of the limits of up to an order of magnitude compared to the 2016-only analysis. The least improvement is found in the region of 250-400 GeV of the electron neutrino coupling scenario where a slight signal-like excess was observed in the same-sign electron + muon channels. Here, we also find the largest divergence between the expected and observed limit. Above 1.2 TeV (1.5 TeV), the observed limit on the coupling squared for the electron neutrino coupling (muon neutrino coupling) crosses 1. As coupling squared values higher than one are not physical, these mass values are also where the exclusion power of the analysis ends. In the case of exclusive muon neutrino coupling, we additionally make the comparison to Ref. [102], which probes this coupling type in the t-channel. It can be seen that it outperforms our search in the high-mass region from 700 GeV onwards.

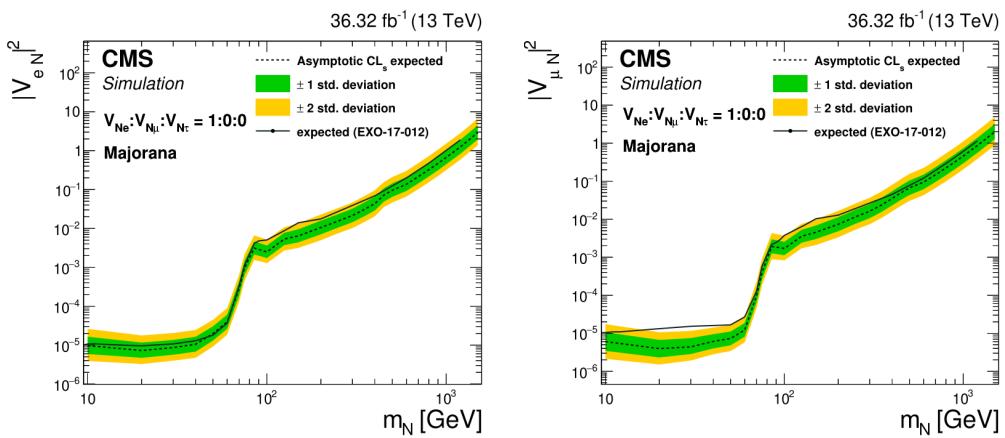


Figure 7.39: Expected 95% CL limits on $|V_{eN}|^2$ (left) and $|V_{\mu N}|^2$ (right) as a function of HNL mass in the assumption of a prompt Majorana HNL using the 2016 dataset. The results are compared to the observed limits from the previous iteration of this analysis [99] (black line).

Interpretations of the results were not only performed for the Majorana assumption but also for the Dirac hypothesis. The limits that roll out of this study are shown in Figure 7.41. One can see that due to the loss of lepton number violating events and thus loss of sensitivity in the search channels without OSSF, the expected limits for the Dirac scenarios are less stringent and are found on average a factor 2 higher than the Majorana limits. One interesting point to note is that in the mass region between 250-400 GeV, the inverse effect is observed in the electron neutrino coupling scenario compared to the Majorana limit. Because the observed excess was found in the same-sign electron + muon channel (which is now no longer sensitive to the signal) and there actually was a slight overprediction in the other channels, the observed limit for a Dirac HNL is found to be lower than the expected limit.

7.10.3 Exclusive tau neutrino coupling

In a similar fashion as in the case of exclusive coupling to light flavour neutrinos, expected limits were produced to study the sensitivity of our search channels. Particularly of interest was the sensitivity contribution of different final states to the final limit. To this end, expected

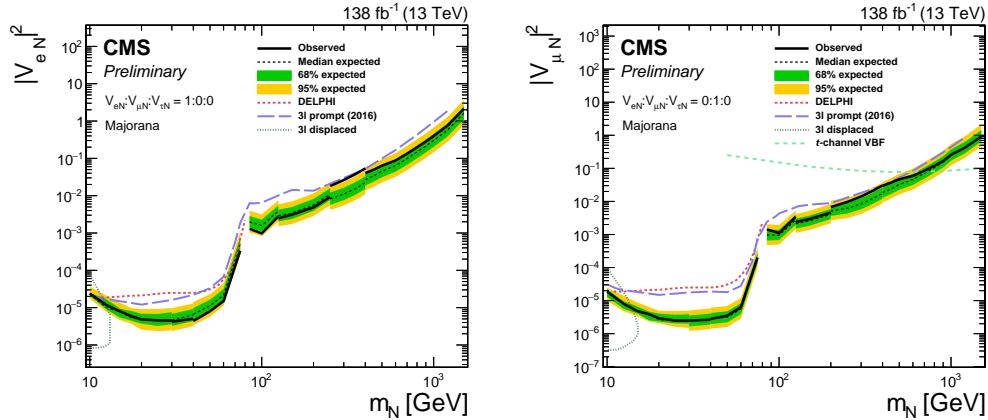


Figure 7.40: Observed 95% CL limits on $|V_{eN}|^2$ (left) and $|V_{\mu N}|^2$ (right) as a function of HNL mass in the assumption of a Majorana HNL using Run II data. The results are compared to the observed limits from the previous iteration of this analysis [99] (blue line), DELPHI [95] (red line) and a displaced HNL search by CMS [96] (dark green line). In the muon coupling scenario, the observed limit from Ref. [102] is shown in light green for comparison.

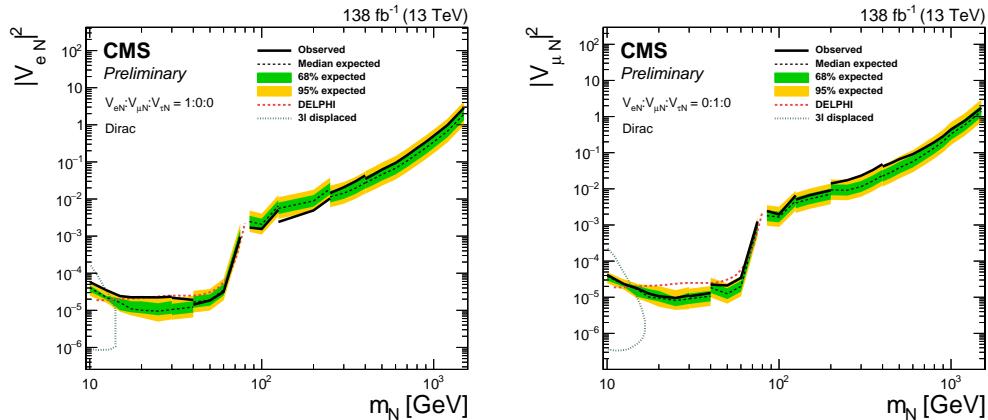


Figure 7.41: Observed 95% CL limits on $|V_{eN}|^2$ (left) and $|V_{\mu N}|^2$ (right) as a function of HNL mass in the assumption of a Dirac HNL using Run II data. The results are compared to the observed limits from DELPHI [95] (red line) and a displaced HNL search by CMS [96] (dark green line).

limits using only light lepton final states, only final states with a single τ_h and two light leptons and the combination of these two final states were produced. They are displayed in Figure 7.42. It is interesting to see that in the low-mass region, the sensitivity is dominated by the light lepton final states. This is not unexpected as the phase space available to the final state leptons is limited and the p_T requirement for τ_h identification is relatively high at 20 GeV. When going to high masses, the roles are reversed and the final states with a single τ_h and two light leptons take over the dominating role in the total sensitivity.

The first observed limits on exclusive τ coupling from CMS are shown in Figure 7.43. In the low-mass part of this plot, we can see exclusion of squared coupling values of 9.0×10^{-4} at 20 GeV going up to 2.5×10^{-3} for a mass point of 60 GeV and reaching 3.3×10^{-2} when

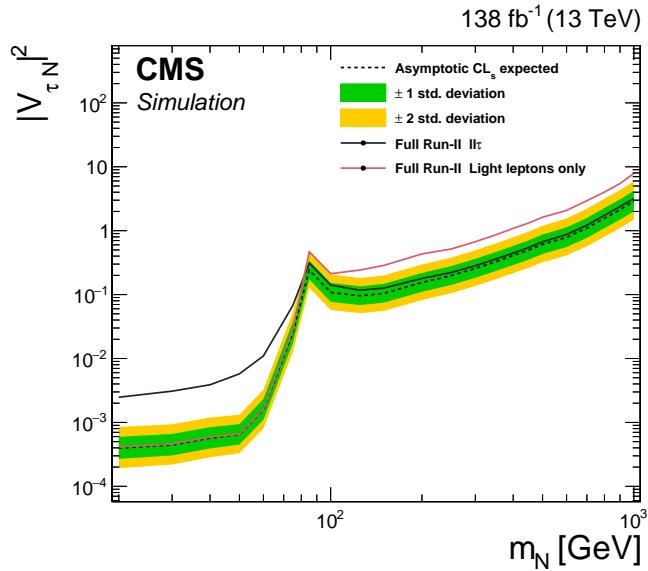


Figure 7.42: Expected 95% CL limits on $|V_{\tau N}|^2$ as a function of HNL mass in the assumption of a prompt Majorana HNL using Run II. The dotted line shows the limit obtained by combining all final states, the full black line the limit obtained by using only final states with a single τ_h and two light leptons and the full red line the limit obtained using fully leptonic final states.

reaching the W mass corridor at 75 GeV where the available phase space for the decay products is limited. When crossing this corridor to the high-mass part we observe exclusion values of 1.5×10^{-1} at 85 GeV, a slight drop to 5.4×10^{-2} around 100 GeV and then a steady climb to 1 at 800 GeV where we lose exclusion power.

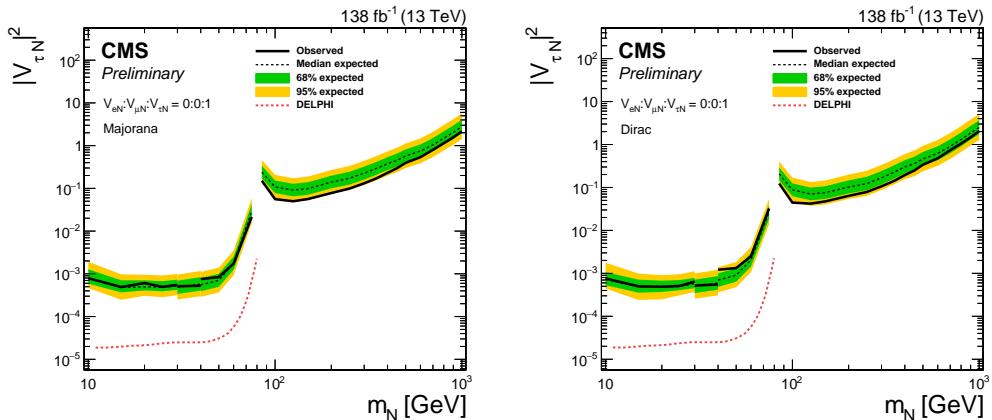


Figure 7.43: Observed 95% CL limits on $|V_{\tau N}|^2$ as a function of HNL mass in the assumption of a Majorana HNL (left) and a Dirac HNL (right) using Run II data. The results are compared to the observed limits from DELPHI [95] (red line).

7.11 Summary and outlook

This chapter presented a search for HNLs in events with three prompt, charged leptons. It aims to shine a light on the presence of heavy neutrinos in the mass range between 10 GeV and 1.5 TeV using benchmark models with exclusive coupling of a single HNL to one SM neutrino flavour. While it is not yet published, it is in the final stages of internal review and subsequent submission to a journal for publication.

Two major mass ranges are taken into account and are studied orthogonally in this analysis. The first is the low-mass region where HNL masses lower than the mass of the W boson are considered. This region is characterized by an off-shell W from HNL decay, resulting in a soft final state lepton spectrum. The second is the high-mass region where HNL masses are higher than the mass of the W boson. Here, the off-shell nature of the initial W boson allows for a less limited phase space. For both regions, final states allowing up to one hadronically decayed τ lepton are studied using a combination of a cut-based strategy of search regions and the use of machine learning in the form of BDTs. The final results show no statistically significant excess of events over the predicted SM background. Finally, 95% CL upper limits were set on the mixing parameters $|V_{eN}|^2$, $|V_{\mu N}|^2$ and for the first time from CMS also $|V_{\tau N}|^2$ as a function of the mass in both the Majorana and Dirac hypothesis. The excluded couplings range from 2.8×10^{-6} to 2.1 in the light flavour mixing scenarios in the aforementioned mass range while the excluded couplings range from 7.7×10^{-4} to 2.1 for tau neutrino mixing scenarios in a mass range from 10 GeV to 1 TeV.

The models considered in this chapter are benchmark models where only one element of the coupling matrix is non-zero, meaning there is only coupling between the HNL and a single flavour. This is done to improve the design of a search that highlights different aspects of HNL models. In reality, these benchmark models are too simplistic and it is unrealistic that the HNL exclusively couples to a single SM neutrino flavour. For this reason, extra care was taken that we provide all necessary information needed for reinterpretation analyses by external parties. To that end, preliminary tests of such reinterpretation analyses are being performed at UGent as part of a master thesis project. As such this analysis serves as a foundation for phenomenologists to perform reinterpretation of the results with the model of their choice.

As this search exclusively targets signatures of promptly decaying HNLs, it runs into limitations at the lowest regarded masses. As the mean lifetime of HNLs is proportional to $m_N^{-5} V_{\ell N}^{-2}$, for masses below 20 GeV we start to see events with non-negligible lifetime on the scale of CMS event propagation, which results in displaced leptons. The prompt impact parameter requirements on the final state leptons remove these events from final selection and effectively lower our acceptance of the signal models at the lowest masses, an effect which increases with lower HNL mass. Two searches performed within the UGent group aim to fill this gap by targeting displaced signatures of low mass HNLs.

The first targets the same Feynman diagram and trilepton final state as this analysis but with the two leptons resulting from HNL decay coming from a secondary vertex and was performed

by Martina Vit, a former PhD student at Ghent University [212]. It was published in Ref. [96]. The search focuses on light lepton final states and couplings between an HNL and light flavour SM neutrinos. Also there, no statistically significant deviation from the predicted SM background is observed. The analysis excludes couplings in the range 3×10^{-7} to 1×10^{-3} for a mass range of 1 to 15 GeV. The second search is performed by Basile Vermassen, an ongoing PhD, and focuses on the decay channel where the W boson from HNL decay proceeds to decay hadronically. This ends up in a final state with a prompt lepton, a displaced lepton, a displaced jet and a secondary displaced vertex. The search will use a machine learning technique called Particle Flow Network (PFN) [213] to improve the separation between signal and background. It is not yet published but will provide exclusion limits in a different channel as this analysis.

The results in this chapter for scenarios with exclusive coupling to electron neutrinos or muon neutrinos improve the existing results by CMS and ATLAS [98, 99], where the same final state is targetted using the 2016 dataset, by an order of magnitude. At the time of writing, this makes them the strongest limits on Type-I Majorana and Dirac neutrinos with this coupling scheme in the m_N range of 30-700 GeV. At lower masses, displaced searches from CMS and ATLAS [96, 97] provide more stringent limits while for $m_N > 700$ GeV the specialized searches for VBF processes by CMS and ATLAS [102, 103] provide better performance. In terms of coupling schemes with exclusive coupling to τ neutrinos, these results are the first direct, prompt limits at the LHC. The only other results related to τ neutrino coupling consist of a mixed coupling interpretation of the displaced search results in Ref. [97] and a CMS search targetting models with a final state of opposite-sign light leptons of which one is displaced in combination with displaced jets [100].

It should be mentioned that within CMS, people are working on an HNL summary paper that lists all variations of HNL research performed within CMS. It is still a work in progress but there are plans to do a combined interpretation of a few different prompt Type-I searches performed using CMS data as part of this effort. These searches look for similar HNLs but produced in different channels: the three charged lepton final state (this search), the two same-sign lepton + jet final state [101] and the VBF channel [102].

Finally, we can take a step back and use what we have learnt to see where the analysis could be improved if there was more time, resources and/or person power. One such item concerns the machine learning strategy. While the BDTs perform very well, they cause discontinuities in the final combination. Because the kinematics between different mass point ranges vary quite rapidly, a series of BDTs needed to be trained and interpreted in order to maintain high discrimination power for all mass points considered. One interesting alternative to this strategy is a parametric neural network. Instead of training a different machine learning algorithm for different kinematics regions, one would have to train just one neural network where the mass of the HNL is a parameter to the model. This network would then be able to efficiently interpolate to other mass points and take into account the varying kinematics. Unfortunately, due to time constraints, this strategy could not be investigated.

Another point where improvements can be made is found in the channels containing one τ_h and two light leptons in the high-mass region. They are included as part of a larger search strategy but could use more specific tweaking. Because the τ_h 's in the high-mass region have a higher p_T spectrum, one could look into using tau triggers in the trigger strategy. Other aspects can be found in the fact that the hadronic decay of a τ lepton contains a neutrino, which means that part of the energy will go to the transverse missing momentum. This could alter the kinematics of the event and the selection strategy could benefit from a more tailored approach. Originally, the τ_h final states in the high-mass region were handled by a different group in close collaboration with our group but due to circumstances, they could not continue on it. Because we observed a high contribution to the sensitivity from these channels even when using the same binning and selection as for the light lepton final states as seen in Figure 7.42, we decided to include these channels as-is. Other groups within CMS are now working on an analysis to reprocess this final state in the high-mass region with a higher sensitivity.

Personal contributions

The analysis presented here is the work of multiple people performed as part of the CMS collaboration. I functioned as the main analyser for it, working on all aspects except for the nonprompt light lepton fake rate estimation and closure tests, which were performed by Luka Lambrecht and the design and validation of the lepton MVA, which was done by Kirill Skovpen. The design of the trigger strategy, the tau ID, updated analysis strategy, nonprompt τ_h estimation, prompt background estimation, systematics and interpretation were performed by me. The documentation of the analysis note and HEPdata output was also written by me. Additionally, I performed a synchronization with the previous iteration of the analysis. On the technical side, I built the analysis framework that runs all aspects of the analysis as well the pythia filters for the signal samples. I have presented all progress in the internal meetings and implemented feedback from my colleagues at UGent and ULB into the analysis. Finally, I have given the pre-approval and approval presentations for the analysis review and presented the analysis at various workshops and CMS meetings.

8

Conclusions

8.1 Summary

This dissertation documents two searches for physics beyond the SM using proton-proton collision data at a centre-of-mass energy of 13 TeV recorded by the CMS detector at the LHC. These searches probe models that aim to address some of the open questions that the SM is not yet able to solve. The first model under investigation, SUSY, reaches this goal by introducing new symmetries and associated new particles. The second model focuses on the neutrino sector of the SM and brings heavy neutral leptons, referred to as HNLs, into the mix. While both models discuss widely different aspects of SM physics, their respective searches share the same final state of interest: signatures in purely leptonic final states. The advantage of this is that in a proton-proton collider, leptons have relatively clean signatures with high efficiencies and low misidentification rates while hadronic final states suffer somewhat more from the high QCD activity in the LHC environment.

The first search, discussed in chapter 6, looks for the presence of electroweak superpartners in the data. It does this using final states containing multiple leptons, starting from at least two leptons. Using an approach that scans a large range of final states, it serves as a general search probing a variety of simplified models. With the use of novel machine learning techniques, improvements of up to a factor 10 in cross section exclusion for the relevant models could be obtained compared to the state-of-the-art at the time. No excess of events pointing towards signs of SUSY was observed in the data. The exclusion limits obtained still serve as the most competitive limits out there, however. I was not the main analyser for this analysis but directly contributed to it by optimising the identification of hadronically decayed τ leptons (τ_h), and developing and validating the nonprompt τ_h background estimation technique.

The second search, for which I was the main analyser and which is presented in chapter 7, looks for the production of HNLs and their subsequent instantaneous decay to two leptons and a SM neutrino at CMS. Final states of exactly three leptons, of which at most one is allowed to be a τ_h , are employed to scan the data for signatures of HNLs with masses in the range of 10 GeV to 1.5 TeV in both the Majorana and Dirac hypotheses. The signal models used are benchmark models with only one HNL that couples exclusively to a single SM neutrino

flavour. No significant excess of events was observed over the SM background prediction and 95% exclusion limits on the mixing parameters $|V_{\ell N}|^2$ as a function of HNL mass were placed. The search improves on the state-of-the-art results in the same production channel by an order of magnitude in scenarios with exclusive coupling to one of the light leptons and is the most competitive limit in the range of 30 to 700 GeV in Type-I HNLs overall for these coupling scenarios. It also provides the first result with sensitivity to scenarios with exclusive coupling between the τ neutrino and an instantaneously decaying HNL.

8.2 Outlook

In this final section of the dissertation, let us take a moment to step into the bigger picture, see where the performed research fits in and what the future holds for the relevant fields. Before we start separating the two searches, let us first look at two limitations they are both affected by, namely the size of the dataset and the centre-of-mass energy of the LHC. An increase in the first improves the statistical precision and significance and gives the power to detect even rarer processes. A higher centre-of-mass energy broadens the range of new particle masses to explore and allows to push the searches to even higher masses.

Currently, Run III is in progress at the LHC. This has a slight centre-of-mass energy bump to 13.6 TeV and is expected to have an integrated luminosity twice as high as Run II. The next big step will come with the High-Luminosity (HL) upgrade of the LHC, which is planned to start at the end of 2027 [214, 215]. It is expected to bring a factor 5 increase in the luminosity and a factor 10 in integrated luminosity compared to Run I to III combined. A potential major upgrade in collision energy at CERN could come with the construction of the Future Circular Collider (FCC). There are currently feasibility studies ongoing for this project. It would consist of a tunnel 100 km in circumference that can reach centre-of-mass energies of 100 GeV. We will sketch the impact of these upgrades on the different searches among other subjects in the following paragraphs.

It is not sure how much the increased dataset will help SUSY searches. At this point, most final states and regular channels have been studied with full Run II data and all of them report no excess of events over the SM background. In fact, some of them have exclusion limits that go to such high sparticle masses that if these sparticles exist and reside at that mass, new hierarchy issues need to be considered. On the other hand, the cross section of superpartner production decreases rapidly when going to higher masses if one stays at the same centre-of-mass energy. This will make it increasingly difficult to keep expanding the existing exclusion limits, even with a large dataset.

Nonetheless, there are some avenues one can go down in order to further SUSY research. Firstly, analysis techniques and strategies keep evolving as technology and ingenuity keep evolving. By playing around with novel techniques such as the parametric neural network employed here, one can keep increasing the potential of the current datasets and datasets to come. Secondly, individual analyses tend to focus on a specific process and are quite

self-contained. That does not mean they do not target similar topologies. By combining the results of these searches, one can get more stringent limits in a more complete phase space. In fact, there are currently ongoing efforts within CMS to perform such a combination that targets electroweak production of winos, binos, higgsinos and sleptons [216]. The results from the analysis described in this dissertation are also included in this combination. Another road to explore is one of exotic topologies that are not currently covered by the current searches. All-in-all, SUSY research will most likely benefit the most from a significant increase in centre-of-mass energy as brought by future colliders.

Searching for additional neutral leptons in the SM is booming business. There are a lot of exclusion results available for the sub-GeV HNL mass region but the higher masses were relatively unexplored until a few years ago. A few analyses at CMS, including the one described in this dissertation, addressed this by exploring the mass range between 10 GeV and 25 TeV. This was done in both the W boson decay channel, with both leptonic and hadronic final states, and in VBF processes. The next step is to perform a combination of these channels and produce combined exclusion limits. This is currently ongoing and the analysis described in chapter 7 will be a part of it. One other interesting final state to study is the one where the second W boson decays hadronically and the first two leptons have the same sign. Up to now, there is one result on this process from CMS where only the 2016 dataset was used. No full Run II results are published yet. As the W has twice the decay rate for hadronic channels, it would be an interesting analysis to update. Especially in the mass range above the W boson mass, the analysis provides more stringent results than the 2016 trilepton analysis.

This analysis was the first one to provide sensitivity to instantaneously decaying HNLs that mix with τ neutrinos. It is a relatively unexplored road in the high-mass range and is something that has been repeatedly requested by phenomenologists. While the initial limits are now in place, there is room for improvement. Within CMS a search is in development that specifically focuses on this coupling scenario in the range of HNL masses above the W boson mass that will continue the pursuit to better understand third-generation coupling to HNLs.

Large improvements are to be expected by the large dataset at the end of the HL-LHC. It will allow to keep pushing the exclusion limits to models with lower coupling squared values and thus lower cross sections. Next-generation particle colliders will keep on improving on these results. This is not limited to hadron colliders, however. If, for example, FCC-ee comes to fruition running up to the proposed FCC experiment – in a similar fashion as LHC being housed in the same tunnel as LEP – one can perform searches for HNLs in the Z-boson decay. The last such search by DELPHI provided limits that are competitive to this day. Projections for FCC-ee show expected exclusion limits down to 10^{-12} in the mass range 10-80 GeV for a range of mixing scenarios [217].

When going to HNL masses in the low-GeV range, searches that target the instantaneous decay of the HNL such as the one described in this dissertation lose sensitivity due to the long-lived nature of the sterile particles. Specialized searches focusing on signatures where the HNL decays outside of the first tracker layers, referred to as displaced, can provide

better performance in that region. CMS and ATLAS have already published multiple such searches and will continue to do so. There are other experiments lined up that will exploit the displaced nature of HNLs in the near or far future. Two of them are FASER and MATHUSLA, which make use of very forward and surface detectors at the LHC interaction points to detect displaced HNLs on a large scale. Also, future beam dump experiments such as SHiP and NA62 show promising prospects. The last word on displaced HNLs has not yet been spoken and we are sure to hear a lot about the topic in the next years.

Last but not least, this analysis does not end at the benchmark models that were researched in this work. As the name states, these models form a benchmark for a specific coupling type. Now that the analysis is designed and unblinded and we do not observe signs of new physics, it can be used to study other coupling scenarios in reinterpretation analyses. Extra effort was taken to provide detailed records of all information needed to perform such analyses. In this context, a pilot study with simple light flavour coupling scenarios was performed at UGent to ensure all needed information is made available. Simultaneously interaction with the phenomenology community is maintained to tailor to their needs. This further leans into the healthy two-way relationship I have always observed between the different groups that I am sure will lead to a prosperous future for the field of HNL.

Bibliography

- [1] D. Riepe, “Naturalistic Tradition in Indian Thought”. Motilal Banarsi Dass Publishers (Pvt. Limited), 1996. ISBN 9788120812932. [Cited on page 1.]
- [2] J. Barnes, “Early Greek Philosophy”. Classics Series. Penguin Books, 1987. ISBN 9780140444612. [Cited on page 1.]
- [3] CMS Collaboration, “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”, *Phys. Lett. B* **716** (2012) 30–61, doi:10.1016/j.physletb.2012.08.021, arXiv:1207.7235. [Cited on page 1.]
- [4] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Phys. Lett. B* **716** (2012) 1–29, doi:10.1016/j.physletb.2012.08.020, arXiv:1207.7214. [Cited on page 1.]
- [5] CMS Collaboration, “A portrait of the Higgs boson by the CMS experiment ten years after the discovery”, *Nature* **607** (2022), no. 7917, 60–68, doi:10.1038/s41586-022-04892-x, arXiv:2207.00043. [Cited on pages 1 and 171.]
- [6] ATLAS Collaboration, “A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery”, *Nature* **607** (2022), no. 7917, 52–59, doi:10.1038/s41586-022-04893-w, arXiv:2207.00092. [Erratum: *Nature* 612, E24 (2022)]. [Cited on page 1.]
- [7] CMS Collaboration, “Search for electroweak production of charginos and neutralinos in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2022) 147, doi:10.1007/JHEP04(2022)147, arXiv:2106.14246. [Cited on pages 2, 78, 79, 82, 89, 108, 109, 110, 112, and 113.]
- [8] CMS Collaboration, “Performance of tau lepton reconstruction at High Level Trigger using 2022 data from the CMS experiment at CERN”, *CDS* (2023). [Cited on pages 3 and 75.]
- [9] H. Jones, “Groups, Representations and Physics”. CRC Press, 2020. ISBN 9781420050295. [Cited on page 5.]
- [10] F. Mandl and G. Shaw, “Quantum Field Theory”. Wiley, Chichester, 2nd edition, 2010. ISBN 978-0-471-49683-0. [Cited on pages 5, 6, and 11.]

- [11] M. Thomson, “Modern particle physics”. Cambridge University Press, New York, 2013. ISBN 9781107034266. [Cited on pages 5, 6, 8, 27, and 38.]
- [12] A. Bettini, “Introduction to Elementary Particle Physics, Second Edition”. Cambridge University Press, 2014. ISBN 9781107779730. [Cited on pages 5 and 6.]
- [13] D. J. Gross, “The role of symmetry in fundamental physics”, *Proceedings of the National Academy of Sciences* **93** (1996), no. 25, 14256–14259, doi:10.1073/pnas.93.25.14256. [Cited on page 5.]
- [14] S. Das and G. Kunstatter, “The central role of symmetry in physics”, arXiv:1609.02038. [Cited on page 5.]
- [15] E. Noether, “Invariant variation problems”, *Transport Theory and Statistical Physics* **1** (jan, 1971) 186–207, doi:10.1080/00411457108231446. [Cited on page 6.]
- [16] P. A. M. Dirac, “On the theory of quantum mechanics”, *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* **112** (1926), no. 762, 661–677. [Cited on page 6.]
- [17] Fermilab, PDG. https://en.wikipedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg. [Cited on page 7.]
- [18] Particle Data Group Collaboration, “Review of Particle Physics”, *PTEP* **2022** (2022) 083C01, doi:10.1093/ptep/ptac097. [Cited on page 9.]
- [19] C. S. Wu et al., “Experimental Test of Parity Conservation in β Decay”, *Phys. Rev.* **105** (1957) 1413–1414, doi:10.1103/PhysRev.105.1413. [Cited on page 9.]
- [20] Q. Shafi and C. Wetterich, “Parity Violation in SLAC Experiment and Left-right Symmetric Gauge Models”, *Phys. Lett. B* **79** (1978) 291–293, doi:10.1016/0370-2693(78)90245-9. [Cited on page 9.]
- [21] E. C. G. Sudarshan and R. e. Marshak, “Chirality invariance and the universal Fermi interaction”, *Phys. Rev.* **109** (1958) 1860–1860, doi:10.1103/PhysRev.109.1860.2. [Cited on page 9.]
- [22] R. P. Feynman and M. Gell-Mann, “Theory of Fermi interaction”, *Phys. Rev.* **109** (1958) 193–198, doi:10.1103/PhysRev.109.193. [Cited on page 9.]
- [23] Particle Data Group Collaboration, “Review of Particle Physics”, *PTEP* **2020** (2020), no. 8, 083C01, doi:10.1093/ptep/ptaa104. [Cited on pages 10, 16, 22, 37, 46, 52, and 54.]
- [24] KATRIN Collaboration, “Direct neutrino-mass measurement with sub-electronvolt sensitivity”, *Nature Phys.* **18** (2022), no. 2, 160–166, doi:10.1038/s41567-021-01463-1, arXiv:2105.08533. [Cited on page 10.]
- [25] E. P. Hincks and B. Pontecorvo, “Search for gamma-radiation in the 2.2-microsecond meson decay process”, *Phys. Rev.* **73** (1948) 257–258, doi:10.1103/PhysRev.73.257. [Cited on page 10.]

- [26] E. J. Konopinski and H. M. Mahmoud, “The Universal Fermi interaction”, *Phys. Rev.* **92** (1953) 1045–1049, doi:10.1103/PhysRev.92.1045. [Cited on page 10.]
- [27] H. Goldstein, C. Poole, and J. Safko, “Classical Mechanics”. Pearson, 2013. ISBN 9781292026558. [Cited on page 11.]
- [28] M. Dasgupta, “An introduction to quantum field theory”, 2008. https://www.ppd.stfc.ac.uk/Pages/Dasgupta_08_Intro_to_QFT.pdf. [Cited on page 11.]
- [29] A. J. Beekman, L. Rademaker, and J. van Wezel, “An Introduction to Spontaneous Symmetry Breaking”, *SciPost Phys. Lect. Notes* **11** (2019) 1, doi:10.21468/SciPostPhysLectNotes.11, arXiv:1909.01820. [Cited on page 13.]
- [30] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons”, *Phys. Rev. Lett.* **13** (1964) 321–323, doi:10.1103/PhysRevLett.13.321. [Cited on page 13.]
- [31] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons”, *Phys. Rev. Lett.* **13** (1964) 508–509, doi:10.1103/PhysRevLett.13.508. [Cited on page 13.]
- [32] T. W. B. Kibble, “Spontaneous symmetry breaking in gauge theories”, *Phil. Trans. Roy. Soc. Lond. A* **373** (2014), no. 2032, 20140033, doi:10.1098/rsta.2014.0033. [Cited on page 13.]
- [33] A. Djouadi, “The Anatomy of electro-weak symmetry breaking. II. The Higgs bosons in the minimal supersymmetric model”, *Phys. Rept.* **459** (2008) 1–241, doi:10.1016/j.physrep.2007.10.005, arXiv:hep-ph/0503173. [Cited on pages 16, 19, and 22.]
- [34] S. Bilenky, “Neutrinos: Majorana or Dirac?”, arXiv:2008.02110. [Cited on page 16.]
- [35] C. Kraus et al., “Final results from phase II of the Mainz neutrino mass search in tritium beta decay”, *Eur. Phys. J. C* **40** (2005) 447–468, doi:10.1140/epjc/s2005-02139-7, arXiv:hep-ex/0412056. [Cited on page 16.]
- [36] Planck Collaboration, “Planck 2015 results. XIII. Cosmological parameters”, *Astron. Astrophys.* **594** (2016) A13, doi:10.1051/0004-6361/201525830, arXiv:1502.01589. [Cited on page 16.]
- [37] KATRIN Collaboration, “Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN”, *Phys. Rev. Lett.* **123** (2019), no. 22, 221802, doi:10.1103/PhysRevLett.123.221802, arXiv:1909.06048. [Cited on page 16.]

[38] A. V. Zasov, A. S. Saburova, A. V. Khoperskov, and S. A. Khoperskov, “Dark matter in galaxies”, *Physics Uspekhi* **60** (April, 2017) 3, doi:10.3367/UFNe.2016.03.037751, arXiv:1710.10630. [Cited on pages 16 and 17.]

[39] G. Bertone and D. Hooper, “History of dark matter”, *Rev. Mod. Phys.* **90** (2018), no. 4, 045002, doi:10.1103/RevModPhys.90.045002, arXiv:1605.04909. [Cited on page 17.]

[40] M. Bauer and T. Plehn, “Yet Another Introduction to Dark Matter: The Particle Physics Approach”, volume 959 of *Lecture Notes in Physics*. Springer, 2019. doi:10.1007/978-3-030-16234-4. [Cited on page 17.]

[41] D. Clowe et al., “A direct empirical proof of the existence of dark matter”, *Astrophys. J. Lett.* **648** (2006) L109–L113, doi:10.1086/508162, arXiv:astro-ph/0608407. [Cited on page 17.]

[42] C. Quigg, “Cosmic Neutrinos”, in *35th SLAC Summer Institute on Particle Physics: Dark matter: From the cosmos to the Laboratory*. 2, 2008. arXiv:0802.0013. [Cited on page 18.]

[43] L. Canetti, M. Drewes, and M. Shaposhnikov, “Matter and Antimatter in the Universe”, *New J. Phys.* **14** (2012) 095012, doi:10.1088/1367-2630/14/9/095012, arXiv:1204.4186. [Cited on page 18.]

[44] A. D. Sakharov, “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe”, *Pisma Zh. Eksp. Teor. Fiz.* **5** (1967) 32–35, doi:10.1070/PU1991v034n05ABEH002497. [Cited on page 18.]

[45] AMS Collaboration, “Search for anti-helium in cosmic rays”, *Phys. Lett. B* **461** (1999) 387–396, doi:10.1016/S0370-2693(99)00874-6, arXiv:hep-ex/0002048. [Cited on page 18.]

[46] M. Kobayashi and T. Maskawa, “CP Violation in the Renormalizable Theory of Weak Interaction”, *Prog. Theor. Phys.* **49** (1973) 652–657, doi:10.1143/PTP.49.652. [Cited on page 18.]

[47] A. Bilal, “Introduction to supersymmetry”, arXiv:hep-th/0101055. [Cited on page 19.]

[48] S. P. Martin, “A Supersymmetry primer”, *Adv. Ser. Direct. High Energy Phys.* **18** (1998) 1–98, doi:10.1142/9789812839657_0001, arXiv:hep-ph/9709356. [Cited on pages 19 and 22.]

[49] R. Penco, “An Introduction to Effective Field Theories”, arXiv:2006.16285. [Cited on page 20.]

[50] T. T. Dumitrescu and Z. Komargodski, “Aspects of supersymmetry and its breaking”, *Nucl. Phys. B Proc. Suppl.* **216** (2011) 44–68, doi:10.1016/j.nuclphysbps.2011.04.149. [Cited on page 20.]

- [51] C. Csaki, “The Minimal supersymmetric standard model (MSSM)”, *Mod. Phys. Lett. A* **11** (1996) 599, doi:10.1142/S021773239600062X, arXiv:hep-ph/9606414. [Cited on page 21.]
- [52] J. Alwall, P. Schuster, and N. Toro, “Simplified Models for a First Characterization of New Physics at the LHC”, *Phys. Rev. D* **79** (2009) 075020, doi:10.1103/PhysRevD.79.075020, arXiv:0810.3921. [Cited on pages 23 and 79.]
- [53] CMS Collaboration, “Interpretation of Searches for Supersymmetry with Simplified Models”, *Phys. Rev. D* **88** (2013), no. 5, 052017, doi:10.1103/PhysRevD.88.052017, arXiv:1301.2175. [Cited on pages 23 and 79.]
- [54] LHC New Physics Working Group Collaboration, “Simplified Models for LHC New Physics Searches”, *J. Phys. G* **39** (2012) 105005, doi:10.1088/0954-3899/39/10/105005, arXiv:1105.2838. [Cited on pages 23 and 79.]
- [55] L. Wezenbeek, “A study of the production mass as a discriminating variable in a supersymmetry search at the LHC”, Master’s thesis, Ghent University, Ghent, 2018. [Cited on pages 23 and 77.]
- [56] CMS Collaboration, “Lhc susy cross sections”, 2021. <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections>. [Cited on page 23.]
- [57] ATLAS Collaboration, “SUSY Summary Plots March 2022”, technical report, CERN, Geneva, 2022. [Cited on page 24.]
- [58] CMS Collaboration, “Cms susy physics results”, 2023. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS>. [Cited on pages 25 and 26.]
- [59] P. B. Pal, “Dirac, Majorana and Weyl fermions”, *Am. J. Phys.* **79** (2011) 485–498, doi:10.1119/1.3549729, arXiv:1006.1718. [Cited on page 27.]
- [60] Miranda-Sanchez, Lopez-Bonilla and Vazquez-Alvarez, “Dirac, weyl and majorana representations of the gamma matrices”, *IDOSI* (2021) doi:10.5829/idosi.sns.2021.24.28. [Cited on page 27.]
- [61] E. Majorana and L. Maiani, “A symmetric theory of electrons and positrons”, pp. 201–233. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. doi:10.1007/978-3-540-48095-2_10. [Cited on page 27.]
- [62] S. Willenbrock, “Symmetries of the standard model”, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Physics in $D \geq 4$* , pp. 3–38. 10, 2004. arXiv:hep-ph/0410370. [Cited on page 28.]
- [63] P. Hernandez, “Neutrino Physics”, in *8th CERN–Latin-American School of High-Energy Physics*, pp. 85–142. 2016. arXiv:1708.01046. doi:10.5170/CERN-2016-005.85. [Cited on page 29.]

- [64] M. Magg and C. Wetterich, “Neutrino Mass Problem and Gauge Hierarchy”, *Phys. Lett. B* **94** (1980) 61–64, doi:10.1016/0370-2693(80)90825-4. [Cited on page 29.]
- [65] R. N. Mohapatra and G. Senjanovic, “Neutrino Mass and Spontaneous Parity Nonconservation”, *Phys. Rev. Lett.* **44** (1980) 912, doi:10.1103/PhysRevLett.44.912. [Cited on page 29.]
- [66] P. Agrawal et al., “Probing the Type-II Seesaw Mechanism through the Production of Higgs Bosons at a Lepton Collider”, *Phys. Rev. D* **98** (2018), no. 1, 015024, doi:10.1103/PhysRevD.98.015024, arXiv:1803.00677. [Cited on page 29.]
- [67] S. Ashanujaman and K. Ghosh, “Revisiting type-II see-saw: present limits and future prospects at LHC”, *JHEP* **03** (2022) 195, doi:10.1007/JHEP03(2022)195, arXiv:2108.10952. [Cited on page 29.]
- [68] E. J. Chun and P. Sharma, “Search for a doubly-charged boson in four lepton final states in type II seesaw”, *Phys. Lett. B* **728** (2014) 256–261, doi:10.1016/j.physletb.2013.11.056, arXiv:1309.6888. [Cited on page 29.]
- [69] E. Ma, “Pathways to naturally small neutrino masses”, *Phys. Rev. Lett.* **81** (1998) 1171–1174, doi:10.1103/PhysRevLett.81.1171, arXiv:hep-ph/9805219. [Cited on page 29.]
- [70] CMS Collaboration, “Inclusive nonresonant multilepton probes of new phenomena at $\sqrt{s}=13$ TeV”, *Phys. Rev. D* **105** (2022), no. 11, 112007, doi:10.1103/PhysRevD.105.112007, arXiv:2202.08676. [Cited on page 29.]
- [71] S. Boyd, “Neutrino lecture writeups”, 2019.
https://warwick.ac.uk/fac/sci/physics/staff/academic/boyd/stuff/neutrinolectures/lec_neutrinomass_writeup.pdf.
[Cited on page 29.]
- [72] K. Moffat, S. Pascoli, and C. Weiland, “Equivalence between massless neutrinos and lepton number conservation in fermionic singlet extensions of the Standard Model”, arXiv:1712.07611. [Cited on page 31.]
- [73] G. G. Ross, “GRAND UNIFIED THEORIES”. 1985. [Cited on page 31.]
- [74] F. F. Deppisch, P. S. Bhupal Dev, and A. Pilaftsis, “Neutrinos and Collider Physics”, *New J. Phys.* **17** (2015), no. 7, 075019, doi:10.1088/1367-2630/17/7/075019, arXiv:1502.06541. [Cited on pages 31 and 117.]
- [75] J. Beacham et al., “Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report”, *J. Phys. G* **47** (2020), no. 1, 010501, doi:10.1088/1361-6471/ab4cd2, arXiv:1901.09966. [Cited on pages 31 and 35.]

- [76] P. D. Bolton, F. F. Deppisch, L. Gráf, and F. Šimkovic, “Two-Neutrino Double Beta Decay with Sterile Neutrinos”, *Phys. Rev. D* **103** (2021), no. 5, 055019, doi:10.1103/PhysRevD.103.055019, arXiv:2011.13387. [Cited on page 32.]
- [77] J. Barea, J. Kotila, and F. Iachello, “Limits on sterile neutrino contributions to neutrinoless double beta decay”, *Phys. Rev. D* **92** (2015) 093001, doi:10.1103/PhysRevD.92.093001, arXiv:1509.01925. [Cited on page 32.]
- [78] Belle Collaboration, “Search for heavy neutrinos at Belle”, *Phys. Rev. D* **87** (2013), no. 7, 071102, doi:10.1103/PhysRevD.87.071102, arXiv:1301.1105. [Erratum: Phys.Rev.D 95, 099903 (2017)]. [Cited on page 32.]
- [79] A. Kobach and S. Dobbs, “Heavy Neutrinos and the Kinematics of Tau Decays”, *Phys. Rev. D* **91** (2015), no. 5, 053006, doi:10.1103/PhysRevD.91.053006, arXiv:1412.4785. [Cited on page 32.]
- [80] G. Bernardi et al., “Further limits on heavy neutrino couplings”, *Phys. Lett. B* **203** (1988) 332–334, doi:10.1016/0370-2693(88)90563-1. [Cited on page 32.]
- [81] NuTeV, E815 Collaboration, “Search for neutral heavy leptons in a high-energy neutrino beam”, *Phys. Rev. Lett.* **83** (1999) 4943–4946, doi:10.1103/PhysRevLett.83.4943, arXiv:hep-ex/9908011. [Cited on page 32.]
- [82] CHARM Collaboration, “A Search for Decays of Heavy Neutrinos in the Mass Range 0.5-GeV to 2.8-GeV”, *Phys. Lett. B* **166** (1986) 473–478, doi:10.1016/0370-2693(86)91601-1. [Cited on page 32.]
- [83] CHARM-II Collaboration, “Precision measurement of electroweak parameters from the scattering of muon-neutrinos on electrons”, *Phys. Lett. B* **335** (1994) 246–252, doi:10.1016/0370-2693(94)91421-4. [Cited on page 32.]
- [84] J. Orloff, A. N. Rozanov, and C. Santoni, “Limits on the mixing of tau neutrino to heavy neutrinos”, *Phys. Lett. B* **550** (2002) 8–15, doi:10.1016/S0370-2693(02)02769-7, arXiv:hep-ph/0208075. [Cited on page 32.]
- [85] DUNE Collaboration, “Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 1: The LBNF and DUNE Projects”, arXiv:1601.05471. [Cited on page 32.]
- [86] LHCb Collaboration, “Search for Majorana neutrinos in $B^- \rightarrow \pi^+ \mu^- \mu^-$ decays”, *Phys. Rev. Lett.* **112** (2014), no. 13, 131802, doi:10.1103/PhysRevLett.112.131802, arXiv:1401.5361. [Cited on page 32.]
- [87] M. Cristinziani, “The SHiP experiment at CERN”, in *3rd World Summit on Exploring the Dark Side of the Universe*, pp. 175–182. 2020. arXiv:2009.06003. [Cited on page 32.]

- [88] SHiP Collaboration, “Sensitivity of the SHiP experiment to Heavy Neutral Leptons”, *JHEP* **04** (2019) 077, doi:10.1007/JHEP04(2019)077, arXiv:1811.00930. [Cited on page 33.]
- [89] FASER Collaboration, “Letter of Intent for FASER: ForwArd Search ExpeRiment at the LHC”, arXiv:1811.10243. [Cited on page 33.]
- [90] FASER Collaboration, “Technical Proposal for FASER: ForwArd Search ExpeRiment at the LHC”, arXiv:1812.09139. [Cited on page 33.]
- [91] MATHUSLA Collaboration, “Explore the lifetime frontier with MATHUSLA”, *JINST* **15** (2020), no. 06, C06026, doi:10.1088/1748-0221/15/06/C06026, arXiv:1901.04040. [Cited on page 33.]
- [92] CERN, “Faser”, 2023.
<https://home.cern/science/experiments/faser>. [Cited on page 33.]
- [93] J. L. Feng et al., “The Forward Physics Facility at the High-Luminosity LHC”, *J. Phys. G* **50** (2023), no. 3, 030501, doi:10.1088/1361-6471/ac865e, arXiv:2203.05090. [Cited on page 33.]
- [94] A. Atre, T. Han, S. Pascoli, and B. Zhang, “The Search for Heavy Majorana Neutrinos”, *JHEP* **05** (2009) 030, doi:10.1088/1126-6708/2009/05/030, arXiv:0901.3589. [Cited on pages 33, 117, 120, and 123.]
- [95] DELPHI Collaboration, “Search for neutral heavy leptons produced in Z decays”, *Z. Phys. C* **74** (1997) 57–71, doi:10.1007/s002880050370. [Erratum: Z.Phys.C 75, 580 (1997)]. [Cited on pages 33, 189, and 190.]
- [96] CMS Collaboration, “Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **07** (2022) 081, doi:10.1007/JHEP07(2022)081, arXiv:2201.05578. [Cited on pages 33, 171, 189, and 192.]
- [97] ATLAS Collaboration, “Search for heavy neutral leptons in decays of W bosons using a dilepton displaced vertex in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”, arXiv:2204.11988. [Cited on pages 33, 34, and 192.]
- [98] ATLAS Collaboration, “Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector”, *JHEP* **10** (2019) 265, doi:10.1007/JHEP10(2019)265, arXiv:1905.09787. [Cited on pages 34 and 192.]
- [99] CMS Collaboration, “Search for heavy neutral leptons in events with three charged leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. Lett.* **120** (2018), no. 22, 221801, doi:10.1103/PhysRevLett.120.221801, arXiv:1802.02965. [Cited on pages 34, 118, 148, 187, 188, 189, and 192.]

- [100] CMS Collaboration, “Search for long-lived heavy neutral leptons with lepton flavour conserving or violating decays to a jet and an electron, muon, or tau lepton”, technical report, CERN, Geneva, 2023. [Cited on pages 34 and 192.]
- [101] CMS Collaboration, “Search for heavy Majorana neutrinos in same-sign dilepton channels in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **01** (2019) 122, doi:10.1007/JHEP01(2019)122, arXiv:1806.10905. [Cited on pages 34 and 192.]
- [102] CMS Collaboration, “Probing heavy Majorana neutrinos and the Weinberg operator through vector boson fusion processes in proton-proton collisions at $\sqrt{s} = 13$ TeV”, arXiv:2206.08956. [Cited on pages 34, 188, 189, and 192.]
- [103] ATLAS Collaboration, “Search for Majorana neutrinos in same-sign WW scattering events from pp collisions at $\sqrt{s} = 13$ TeV”, arXiv:2305.14931. [Cited on pages 34 and 192.]
- [104] L. Canetti and M. Shaposhnikov, “Baryon Asymmetry of the Universe in the NuMSM”, *JCAP* **09** (2010) 001, doi:10.1088/1475-7516/2010/09/001, arXiv:1006.0133. [Cited on page 34.]
- [105] M. Drewes and B. Garbrecht, “Combining experimental and cosmological constraints on heavy neutrinos”, *Nucl. Phys. B* **921** (2017) 250–315, doi:10.1016/j.nuclphysb.2017.05.001, arXiv:1502.00477. [Cited on page 34.]
- [106] O. Ruchayskiy and A. Ivashko, “Restrictions on the lifetime of sterile neutrinos from primordial nucleosynthesis”, *JCAP* **10** (2012) 014, doi:10.1088/1475-7516/2012/10/014, arXiv:1202.2841. [Cited on page 34.]
- [107] S. Antusch and O. Fischer, “Testing sterile neutrino extensions of the Standard Model at the Circular Electron Positron Collider”, *Int. J. Mod. Phys. A* **30** (2015), no. 23, 1544004, doi:10.1142/S0217751X15440042. [Cited on page 34.]
- [108] E. Mobs, “The CERN accelerator complex in 2019”, *CDS* (2019). General Photo. [Cited on page 38.]
- [109] “LHC Machine”, *JINST* **3** (2008) S08001, doi:10.1088/1748-0221/3/08/S08001. [Cited on page 39.]
- [110] J. Knolle, “Measuring luminosity and the $t\bar{t}Z$ production cross section with the CMS experiment”, Sep, 2020. Presented 25 Jun 2020.
<https://cds.cern.ch/record/2744549>. [Cited on page 39.]
- [111] I. Neutelings, “CMS coordinate system”. https://tikz.net/axis3d_cms/. [Cited on page 41.]
- [112] “Pseudorapidity”. https://tikz.net/axis2d_pseudorapidity/. [Cited on page 41.]

[113] “CMS detector design”.
<https://cms.cern/news/cms-detector-design>. [Cited on page 42.]

[114] CMSPublic Web, “Cms tracker detector performance results”, 2022. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/DPGResultsTRK>.
[Cited on page 44.]

[115] CMS Collaboration, G. L. Bayatian et al., “CMS Physics: Technical Design Report Volume 1: Detector Performance and Software”. Technical design report. CMS. CERN, Geneva, 2006. [Cited on page 45.]

[116] CMS Collaboration, “Calibration of the CMS hadron calorimeters using proton-proton collision data at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020), no. 05, P05002, doi:10.1088/1748-0221/15/05/P05002, arXiv:1910.00079. [Cited on page 45.]

[117] CMS Collaboration, “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018), no. 06, P06015, doi:10.1088/1748-0221/13/06/P06015, arXiv:1804.04528. [Cited on pages 47, 50, and 82.]

[118] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017), no. 10, P10003, doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965. [Cited on pages 48 and 49.]

[119] R. Fruhwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Meth. A* **262** (1987) 444–450, doi:10.1016/0168-9002(87)90887-4. [Cited on page 48.]

[120] R. Fruhwirth, W. Waltenberger, and P. Vanlaer, “Adaptive vertex fitting”, *J. Phys. G* **34** (2007) N343, doi:10.1088/0954-3899/34/12/N01. [Cited on page 48.]

[121] CMS Collaboration, “Technical proposal for the Phase-II upgrade of the Compact Muon Solenoid”, CMS Technical Proposal CERN-LHCC-2015-010, CMS-TDR-15-02, 2015. [Cited on page 49.]

[122] CMS Collaboration, F. Beaudette, “The CMS Particle Flow Algorithm”, in *Proceedings, International Conference on Calorimetry for the High Energy Frontier (CHEF 2013): Paris, France, April 22-25, 2013*, pp. 295–304. 2013. arXiv:1401.8155. [Cited on page 49.]

[123] W. Adam, R. Frühwirth, A. Strandlie, and T. Todorov, “Reconstruction of electrons with the gaussian-sum filter in the CMS tracker at the LHC”, *Journal of Physics G: Nuclear and Particle Physics* **31** (jul, 2005) N9–N20, doi:10.1088/0954-3899/31/9/n01. [Cited on page 51.]

- [124] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_t jet clustering algorithm”, *JHEP* **04** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189. [Cited on pages 51 and 125.]
- [125] CMS Collaboration, “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”, *JINST* **12** (2017), no. 02, P02014, doi:10.1088/1748-0221/12/02/P02014, arXiv:1607.03663. [Cited on page 52.]
- [126] CMS Collaboration Collaboration, “Jet energy scale and resolution performance with 13 TeV data collected by CMS in 2016-2018”, *CDS* (2020). [Cited on page 52.]
- [127] CMS Collaboration, “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV”, *JINST* **13** (2018), no. 05, P05011, doi:10.1088/1748-0221/13/05/P05011, arXiv:1712.07158. [Cited on pages 52, 53, and 87.]
- [128] E. Bols et al., “Jet Flavour Classification Using DeepJet”, *JINST* **15** (2020), no. 12, P12012, doi:10.1088/1748-0221/15/12/P12012, arXiv:2008.10519. [Cited on page 52.]
- [129] CMS Collaboration, “Performance of tau-lepton reconstruction and identification in CMS”, *JINST* **7** (2012) P01001, doi:10.1088/1748-0221/7/01/P01001, arXiv:1109.6034. [Cited on page 53.]
- [130] CMS Collaboration, “Reconstruction and identification of τ lepton decays to hadrons and ν_τ at CMS”, *JINST* **11** (2016), no. 01, P01019, doi:10.1088/1748-0221/11/01/P01019, arXiv:1510.07488. [Cited on page 53.]
- [131] CMS Collaboration, “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018), no. 10, P10005, doi:10.1088/1748-0221/13/10/P10005, arXiv:1809.02816. [Cited on page 53.]
- [132] N. Metropolis and S. Ulam, “The monte carlo method”, *Journal of the American Statistical Association* **44** (1949), no. 247, 335–341, doi:10.1080/01621459.1949.10483310. PMID: 18139350. [Cited on page 56.]
- [133] J. Qiang, “Monte Carlo Simulation Techniques”, in *CAS - CERN Accelerator School: Numerical Methods for Analysis, Design and Modelling of Particle Accelerators*. 6, 2020. arXiv:2006.10506. [Cited on page 56.]
- [134] A. Buckley et al., “General-purpose event generators for LHC physics”, *Phys. Rept.* **504** (2011) 145–233, doi:10.1016/j.physrep.2011.03.005, arXiv:1101.2599. [Cited on page 56.]

[135] E. Bothmann et al., “Accelerating LHC event generation with simplified pilot runs and fast PDFs”, *Eur. Phys. J. C* **82** (2022), no. 12, 1128, doi:10.1140/epjc/s10052-022-11087-1, arXiv:2209.00843. [Cited on page 56.]

[136] F. Maltoni, T. McElmurry, R. Putman, and S. Willenbrock, “Choosing the Factorization Scale in Perturbative QCD”, arXiv:hep-ph/0703156. [Cited on page 56.]

[137] T. Plehn, “LHC Phenomenology for Physics Hunters”, in *Theoretical Advanced Study Institute in Elementary Particle Physics: The Dawn of the LHC Era*, pp. 125–180. 2010. arXiv:0810.2281. doi:10.1142/9789812838360_0003. [Cited on page 56.]

[138] GEANT4 Collaboration, “GEANT4—a simulation toolkit”, *Nucl. Instrum. Meth. A* **506** (2003) 250–303, doi:10.1016/S0168-9002(03)01368-8. [Cited on page 57.]

[139] ATLAS, CMS, LHC Higgs Combination Group Collaboration, “Procedure for the LHC Higgs boson search combination in Summer 2011”,. [Cited on pages 58, 111, and 175.]

[140] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics”, *Eur. Phys. J. C* **71** (2011) 1554, doi:10.1140/epjc/s10052-011-1554-0, arXiv:1007.1727. [Erratum: Eur.Phys.J.C 73, 2501 (2013)]. [Cited on pages 58 and 111.]

[141] CMS Collaboration, “The CMS trigger in Run 2”, *PoS EPS-HEP2017* (2017) 523, doi:10.22323/1.314.0523. [Cited on pages 62 and 63.]

[142] CMS Collaboration, “CMS Technical Design Report for the Level-1 Trigger Upgrade”,. [Cited on page 63.]

[143] CMS Collaboration, “Design and performance of the upgrade of the CMS L1 trigger”, *PoS ICHEP2018* (2019) 198, doi:10.22323/1.340.0198. [Cited on page 63.]

[144] CMS Collaboration, “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020), no. 10, P10017, doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165. [Cited on page 63.]

[145] CMS Collaboration, “CMS Run 2 High Level Trigger Performance”, *PoS EPS-HEP2019* (2020) 165, doi:10.22323/1.364.0165. [Cited on page 63.]

[146] CMS Collaboration, “CMS Trigger Performance”, *EPJ Web Conf.* **182** (2018) 02037, doi:10.1051/epjconf/201818202037. [Cited on page 63.]

[147] CMS Collaboration, “Tracking Performance in the CMS High Level Trigger - June 2018”, *CDS* (2018). [Cited on pages 64 and 65.]

[148] CMS Collaboration, “2018 Tau Trigger Reconstruction Comparison”, *CDS* (2018). [Cited on page 67.]

- [149] CMS Collaboration, “New trigger strategies for CMS during Run 3”, technical report, CERN, Geneva, 2022. [Cited on page 67.]
- [150] M. Dordevic and on behalf of the CMS Collaboration, “The cms trigger system”, *Journal of Physics: Conference Series* **2375** (nov, 2022) 012003, doi:10.1088/1742-6596/2375/1/012003. [Cited on pages 67 and 68.]
- [151] A. Bocci et al., “Heterogeneous Reconstruction of Tracks and Primary Vertices With the CMS Pixel Tracker”, *Front. Big Data* **3** (2020) 601728, doi:10.3389/fdata.2020.601728, arXiv:2008.13461. [Cited on page 67.]
- [152] CMS Collaboration, “Performance of Run-3 HLT Track Reconstruction”, *CDS* (2022). [Cited on page 68.]
- [153] K. O’Shea and R. Nash, “An Introduction to Convolutional Neural Networks”, *arXiv e-prints* (November, 2015) arXiv:1511.08458, arXiv:1511.08458. [Cited on page 68.]
- [154] CMS Collaboration, “Identification of hadronic tau lepton decays using a deep neural network”, arXiv:2201.08458. [Cited on pages 69 and 71.]
- [155] CMS Collaboration, “The Phase-2 Upgrade of the CMS Level-1 Trigger”, *CDS* (2020). [Cited on page 74.]
- [156] S. Frixione and B. R. Webber, “Matching NLO QCD computations and parton shower simulations”, *JHEP* **06** (2002) 029, doi:10.1088/1126-6708/2002/06/029, arXiv:hep-ph/0204244. [Cited on pages 80 and 125.]
- [157] J. Alwall et al., “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, *JHEP* **07** (2014) 079, doi:10.1007/JHEP07(2014)079, arXiv:1405.0301. [Cited on pages 80, 120, and 125.]
- [158] P. Nason, “A New method for combining NLO QCD with shower Monte Carlo algorithms”, *JHEP* **11** (2004) 040, doi:10.1088/1126-6708/2004/11/040, arXiv:hep-ph/0409146. [Cited on pages 81 and 125.]
- [159] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with Parton Shower simulations: the POWHEG method”, *JHEP* **11** (2007) 070, doi:10.1088/1126-6708/2007/11/070, arXiv:0709.2092. [Cited on pages 81 and 125.]
- [160] T. Melia, P. Nason, R. Rontsch, and G. Zanderighi, “W+W-, WZ and ZZ production in the POWHEG BOX”, *JHEP* **11** (2011) 078, doi:10.1007/JHEP11(2011)078, arXiv:1107.5051. [Cited on pages 81 and 125.]
- [161] P. Nason and G. Zanderighi, “ W^+W^- , WZ and ZZ production in the POWHEG-BOX-V2”, *Eur. Phys. J. C* **74** (2014), no. 1, 2702,

doi:10.1140/epjc/s10052-013-2702-5, arXiv:1311.1365. [Cited on pages 81 and 125.]

[162] J. M. Campbell and R. K. Ellis, “MCFM for the Tevatron and the LHC”, *Nucl. Phys. B Proc. Suppl.* **205-206** (2010) 10–15, doi:10.1016/j.nuclphysbps.2010.08.011, arXiv:1007.3492. [Cited on pages 81 and 125.]

[163] CMS Collaboration, “Pileup mitigation at CMS in 13 TeV data”, *JINST* **15** (2020), no. 09, P09018, doi:10.1088/1748-0221/15/09/P09018, arXiv:2003.00503. [Cited on pages 82 and 105.]

[164] CMS Collaboration, “Evidence for associated production of a Higgs boson with a top quark pair in final states with electrons, muons, and hadronically decaying τ leptons at $\sqrt{s} = 13$ TeV”, *JHEP* **08** (2018) 066, doi:10.1007/JHEP08(2018)066, arXiv:1803.05485. [Cited on page 82.]

[165] CMS Collaboration, “Observation of Single Top Quark Production in Association with a Z Boson in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. Lett.* **122** (2019), no. 13, 132003, doi:10.1103/PhysRevLett.122.132003, arXiv:1812.05900. [Cited on page 82.]

[166] CMS Collaboration Collaboration, “Performance of b tagging algorithms in proton-proton collisions at 13 TeV with Phase 1 CMS detector”, *CDS* (2018). [Cited on page 82.]

[167] A. Bodek et al., “Extracting Muon Momentum Scale Corrections for Hadron Collider Experiments”, *Eur. Phys. J. C* **72** (2012) 2194, doi:10.1140/epjc/s10052-012-2194-8, arXiv:1208.3710. [Cited on pages 83 and 126.]

[168] CMS Collaboration, “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”, *JINST* **16** (2021), no. 05, P05014, doi:10.1088/1748-0221/16/05/P05014, arXiv:2012.06888. [Cited on pages 84 and 124.]

[169] CMS Collaboration Collaboration, “Jet Performance in pp Collisions at 7 TeV”, technical report, CERN, Geneva, 2010. [Cited on page 87.]

[170] C. G. Lester and D. J. Summers, “Measuring masses of semiinvisibly decaying particles pair produced at hadron colliders”, *Phys. Lett. B* **463** (1999) 99–103, doi:10.1016/S0370-2693(99)00945-4, arXiv:hep-ph/9906349. [Cited on page 90.]

[171] CMS Collaboration, “Search for new physics with same-sign isolated dilepton events with jets and missing transverse energy at the LHC”, *JHEP* **06** (2011) 077, doi:10.1007/JHEP06(2011)077, arXiv:1104.3168. [Cited on page 93.]

[172] CMS Collaboration, “Search for a Standard Model Higgs Boson Produced in Association with a Top-Quark Pair and Decaying to Bottom Quarks Using a Matrix Element Method”, *Eur. Phys. J. C* **75** (2015), no. 6, 251, doi:10.1140/epjc/s10052-015-3454-1, arXiv:1502.02485. [Cited on page 94.]

[173] 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. [Cited on pages 106 and 171.]

[174] T. Junk, “Confidence level computation for combining searches with small statistics”, *Nucl. Instrum. Meth. A* **434** (1999) 435–443, doi:10.1016/S0168-9002(99)00498-2, arXiv:hep-ex/9902006. [Cited on page 111.]

[175] A. L. Read, “Presentation of search results: The CL(s) technique”, *J. Phys. G* **28** (2002) 2693–2704, doi:10.1088/0954-3899/28/10/313. [Cited on page 111.]

[176] J. S. Conway, “Incorporating Nuisance Parameters in Likelihoods for Multisource Spectra”, in *PHYSTAT 2011*, pp. 115–120. 2011. arXiv:1103.0354. doi:10.5170/CERN-2011-006.115. [Cited on page 111.]

[177] R. J. Barlow and C. Beeston, “Fitting using finite Monte Carlo samples”, *Comput. Phys. Commun.* **77** (1993) 219–228, doi:10.1016/0010-4655(93)90005-W. [Cited on page 111.]

[178] CMS Collaboration, “Search for electroweak production of charginos and neutralinos in multilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **03** (2018) 166, doi:10.1007/JHEP03(2018)166, arXiv:1709.05406. [Cited on pages 111, 112, and 113.]

[179] F. del Aguila and J. A. Aguilar-Saavedra, “Distinguishing seesaw models at LHC with multi-lepton signals”, *Nucl. Phys. B* **813** (2009) 22–90, doi:10.1016/j.nuclphysb.2008.12.029, arXiv:0808.2468. [Cited on page 117.]

[180] V. Tello et al., “Left-Right Symmetry: from LHC to Neutrinoless Double Beta Decay”, *Phys. Rev. Lett.* **106** (2011) 151801, doi:10.1103/PhysRevLett.106.151801, arXiv:1011.3522. [Cited on page 117.]

[181] Y. Cai, T. Han, T. Li, and R. Ruiz, “Lepton Number Violation: Seesaw Models and Their Collider Tests”, *Front. in Phys.* **6** (2018) 40, doi:10.3389/fphy.2018.00040, arXiv:1711.02180. [Cited on page 117.]

[182] S. Pascoli, R. Ruiz, and C. Weiland, “Heavy neutrinos with dynamic jet vetoes: multilepton searches at $\sqrt{s} = 14, 27$, and 100 TeV”, *JHEP* **06** (2019) 049, doi:10.1007/JHEP06(2019)049, arXiv:1812.08750. [Cited on pages 117 and 120.]

[183] A. M. Abdullahi et al., “The present and future status of heavy neutral leptons”, *J. Phys. G* **50** (2023), no. 2, 020501, doi:10.1088/1361-6471/ac98f9, arXiv:2203.08039. [Cited on page 117.]

[184] D. Alva, T. Han, and R. Ruiz, “Heavy Majorana neutrinos from $W\gamma$ fusion at hadron colliders”, *JHEP* **02** (2015) 072, doi:10.1007/JHEP02(2015)072, arXiv:1411.7305. [Cited on pages 118, 119, and 120.]

[185] C. Degrande, O. Mattelaer, R. Ruiz, and J. Turner, “Fully-Automated Precision Predictions for Heavy Neutrino Production Mechanisms at Hadron Colliders”, *Phys. Rev. D* **94** (2016), no. 5, 053002, doi:10.1103/PhysRevD.94.053002, arXiv:1602.06957. [Cited on pages 118 and 120.]

[186] M. Drewes, J. Klarić, and J. López-Pavón, “New Benchmark Models for Heavy Neutral Lepton Searches”, arXiv:2207.02742. [Cited on page 119.]

[187] A. Abada, P. Escribano, X. Marcano, and G. Piazza, “Collider searches for heavy neutral leptons: beyond simplified scenarios”, *Eur. Phys. J. C* **82** (2022), no. 11, 1030, doi:10.1140/epjc/s10052-022-11011-7, arXiv:2208.13882. [Cited on page 119.]

[188] R. Ruiz et al., “HeavyN: The standard model + heavy neutrinos at NLO in QCD”, 2021. <https://feynrules.irmpl.ucl.ac.be/wiki/HeavyN>. [Cited on pages 120 and 121.]

[189] NNPDF Collaboration, “Parton distributions from high-precision collider data”, *Eur. Phys. J. C* **77** (2017), no. 10, 663, doi:10.1140/epjc/s10052-017-5199-5, arXiv:1706.00428. [Cited on pages 121 and 125.]

[190] T. Sjöstrand et al., “An introduction to PYTHIA 8.2”, *Comput. Phys. Commun.* **191** (2015) 159–177, doi:10.1016/j.cpc.2015.01.024, arXiv:1410.3012. [Cited on pages 121 and 125.]

[191] CMS Collaboration, “Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements”, *Eur. Phys. J. C* **80** (2020), no. 1, 4, doi:10.1140/epjc/s10052-019-7499-4, arXiv:1903.12179. [Cited on pages 121 and 125.]

[192] J. Alwall et al., “Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions”, *Eur. Phys. J. C* **53** (2008) 473–500, doi:10.1140/epjc/s10052-007-0490-5, arXiv:0706.2569. [Cited on pages 121 and 125.]

[193] R. Frederix and S. Frixione, “Merging meets matching in MC@NLO”, *JHEP* **12** (2012) 061, doi:10.1007/JHEP12(2012)061, arXiv:1209.6215. [Cited on pages 121 and 125.]

- [194] T. Han and B. Zhang, “Signatures for Majorana neutrinos at hadron colliders”, *Phys. Rev. Lett.* **97** (2006) 171804, doi:10.1103/PhysRevLett.97.171804, arXiv:hep-ph/0604064. [Cited on pages 123 and 124.]
- [195] V. Cirigliano et al., “Leptonic anomalous magnetic moments in ν SMEFT”, *JHEP* **08** (2021) 103, doi:10.1007/JHEP08(2021)103, arXiv:2105.11462. [Cited on page 123.]
- [196] M. Jansova, “Search for the supersymmetric partner of the top quark and measurements of cluster properties in the silicon strip tracker of the CMS experiment at Run 2”, 2018. Presented 27 Sep 2018.
<https://cds.cern.ch/record/2647308>. [Cited on page 124.]
- [197] Cavallari, Francesca and Rovelli, Chiara, “Calibration and performance of the cms electromagnetic calorimeter in lhc run2”, *EPJ Web Conf.* **245** (2020) 02027, doi:10.1051/epjconf/202024502027. [Cited on page 124.]
- [198] CMS Collaboration, “ECAL 2016 refined calibration and Run2 summary plots”,. [Cited on page 124.]
- [199] CMS Collaboration, “Observation of four top quark production in proton-proton collisions at $\sqrt{s} = 13$ TeV”, arXiv:2305.13439. [Cited on pages 126, 169, and 171.]
- [200] A. Hoecker et al., “TMVA - Toolkit for Multivariate Data Analysis”, *arXiv e-prints* (March, 2007) physics/0703039, arXiv:physics/0703039. [Cited on pages 126 and 147.]
- [201] Y. Coadou, “Boosted decision trees”, doi:10.1142/9789811234033_0002, arXiv:2206.09645. [Cited on page 145.]
- [202] Y. Freund and R. E. Schapire, “A desicion-theoretic generalization of on-line learning and an application to boosting”, in *Computational Learning Theory*, P. Vitányi, ed., pp. 23–37. Springer Berlin Heidelberg, Berlin, Heidelberg, 1995. [Cited on page 146.]
- [203] CMS Collaboration, “Measurement of the cross section of top quark-antiquark pair production in association with a W boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”, arXiv:2208.06485. [Cited on pages 169 and 171.]
- [204] A. Lazopoulos, K. Melnikov, and F. Petriello, “QCD corrections to tri-boson production”, *Phys. Rev. D* **76** (2007) 014001, doi:10.1103/PhysRevD.76.014001, arXiv:hep-ph/0703273. [Cited on page 170.]
- [205] T. Binoth, G. Ossola, C. G. Papadopoulos, and R. Pittau, “NLO QCD corrections to tri-boson production”, *JHEP* **06** (2008) 082, doi:10.1088/1126-6708/2008/06/082, arXiv:0804.0350. [Cited on page 170.]

[206] F. Campanario et al., “QCD corrections to charged triple vector boson production with leptonic decay”, *Phys. Rev. D* **78** (2008) 094012, doi:10.1103/PhysRevD.78.094012, arXiv:0809.0790. [Cited on page 170.]

[207] S. Dittmaier, A. Huss, and G. Knippen, “Next-to-leading-order QCD and electroweak corrections to WWW production at proton-proton colliders”, *JHEP* **09** (2017) 034, doi:10.1007/JHEP09(2017)034, arXiv:1705.03722. [Cited on page 170.]

[208] CMS Collaboration, “Inclusive and differential cross section measurements of single top quark production in association with a Z boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **02** (2022) 107, doi:10.1007/JHEP02(2022)107, arXiv:2111.02860. [Cited on page 171.]

[209] CMS Collaboration, “Measurement of the inclusive and differential $t\bar{t}\gamma$ cross sections in the dilepton channel and effective field theory interpretation in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **05** (2022) 091, doi:10.1007/JHEP05(2022)091, arXiv:2201.07301. [Cited on page 171.]

[210] M. Cacciari et al., “The t anti-t cross-section at 1.8-TeV and 1.96-TeV: A Study of the systematics due to parton densities and scale dependence”, *JHEP* **04** (2004) 068, doi:10.1088/1126-6708/2004/04/068, arXiv:hep-ph/0303085. [Cited on page 171.]

[211] A. L. Read, “Presentation of search results: the cl s technique”, *Journal of Physics G: Nuclear and Particle Physics* **28** (2002), no. 10, 2693. [Cited on page 175.]

[212] M. Vit, “Search for heavy neutral leptons in pp collision events with three charged leptons using the CMS detector”, Dec, 2021. Presented 14 Dec 2021. <http://lib.ugent.be/catalog/pug01:8734585>. [Cited on page 192.]

[213] P. T. Komiske, E. M. Metodiev, and J. Thaler, “Energy Flow Networks: Deep Sets for Particle Jets”, *JHEP* **01** (2019) 121, doi:10.1007/JHEP01(2019)121, arXiv:1810.05165. [Cited on page 192.]

[214] O. Brüning and L. Rossi, eds., “The High Luminosity Large Hadron Collider: the new machine for illuminating the mysteries of Universe”, volume 24. 2015. doi:10.1142/9581, ISBN 978-981-4675-46-8, 978-981-4678-14-8. [Cited on page 196.]

[215] O. Brüning and L. Rossi, “The High-Luminosity Large Hadron Collider”, *Nature Review Physics* **01** (2019) 241–243, doi:10.1038/s42254-019-0050-6. [Cited on page 196.]

[216] CMS Collaboration, “Combined search for electroweak production of winos, binos, higgsinos, and sleptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, technical report, CERN, Geneva, 2023. [Cited on page 197.]

[217] FCC-ee study Team Collaboration, “Search for Heavy Right Handed Neutrinos at the FCC-ee”, *Nucl. Part. Phys. Proc.* **273-275** (2016) 1883–1890,

doi:10.1016/j.nuclphysbps.2015.09.304, arXiv:1411.5230. [Cited
on page 197.]