

DISS. ETH NO. 30611

*Search for new physics in the
supersymmetric and neutrino sectors*

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES

(Dr. sc. ETH Zürich)

presented by

Anne-Mazarine Lyon

M.Sc. in Applied Physics, École Polytechnique Fédérale de Lausanne

born on 15.07.1995

accepted on the recommendation of

Prof. Dr. Günther Dissertori

Prof. Dr. Jeffrey Richman

2024

Abstract

The success of the standard model (SM) of particle physics in describing the strong, electromagnetic, and weak fundamental interactions is considered as one of the greatest achievements of modern physics. Nevertheless, this theoretical framework is believed to be incomplete, as it leaves unanswered major questions at both the particle and cosmological frontiers. *New physics* refers to possible extensions of the SM that provide solutions to these open questions.

The thesis presents two searches for new physics in the form of new particles predicted in supersymmetric and heavy neutrino models. The presence of these hypothetical particles is sought after in the proton-proton collision data produced at the CERN Large Hadron Collider (LHC) at a centre-of-mass energy of 13 TeV and collected by the CMS experiment.

The first search targets the production of supersymmetric particles giving rise to all-hadronic final states with large transverse momentum imbalance. The analysis is performed using data collected during the Run 2 (2016 – 2018) of the LHC and corresponding to an integrated luminosity of 137 fb^{-1} . Models are considered in which supersymmetric particles are produced via the strong interaction in processes conserving R -parity. The strategy of the search leverages on the high separation power of the M_{T2} variable between events arising from supersymmetric and from SM background processes. The data are found to agree with the SM background predictions. The results are used to constrain at 95% confidence level (CL) the mass parameters of various supersymmetric models. For the majority of these models, the most stringent limits to date are obtained.

The second search for new physics presented in the thesis considers, for the first time in CMS, the production of heavy neutrinos in the decays of beauty mesons. This novel analysis exploits a special beauty-hadron-enriched data sample collected in 2018 that contains of order $10^{10} \text{ b}\bar{\text{b}}$ events for an integrated luminosity of 41.6 fb^{-1} . The heavy neutrinos are assumed to be long lived to study the case in which they mix feebly with the SM neutrinos. No significant excess of events over the SM background is observed. Limits at 95% CL are presented for parameters describing the mass, lifetime, and mixing of the heavy neutrinos. Multiple scenarios are constrained for both the Majorana and Dirac-like hypotheses and for different mixing patterns. Among the constraints obtained, some are the first ever derived and others are the most stringent to date.

Finally, searching for new physics would not be possible without excellent particle reconstruction algorithms at CMS. The last project presented in the thesis consists in the optimisation of one of the reconstruction algorithms used with the electromagnetic calorimeter, performed in preparation for the Run 3 of data taking. All the CMS analyses that exploit Run 3 data and that rely on photons, electrons, jets, or missing transverse momentum benefit from this work.

Résumé

Le succès du modèle standard de la physique des particules dans sa description des interactions forte, électromagnétique, et faible est considéré comme l’une des plus grandes réussites de la physique moderne. Cependant, ce cadre théorique n’est pas complet, étant donné qu’il n’apporte aucune réponse à des questions fondamentales de physique. La *nouvelle physique* fait référence aux possibles extensions du modèle standard qui apportent des solutions à ces questions ouvertes.

La thèse présente deux recherches de nouvelle physique sous la forme de nouvelles particules prédites dans des modèles de supersymétrie et de neutrinos lourds. La présence de ces particules hypothétiques est recherchée dans les données de collisions proton-proton, produites au LHC du CERN à une énergie de centre de masse de 13 TeV et collectées par l’expérience CMS.

La première recherche cible la production de particules supersymétriques qui donnent lieu à un état final ne comportant que des hadrons et une importante énergie manquante. L’analyse est réalisée en utilisant des données collectées pendant le Run 2 (2016 – 2018) du LHC et correspondant à une luminosité intégrée de 137 fb^{-1} . Les modèles scrutés considèrent le cas où les particules supersymétriques sont produites via l’interaction forte dans des processus qui conservent la parité R . La stratégie de l’analyse exploite la capacité de la variable M_{T2} à distinguer les événements résultant de processus supersymétriques de ceux résultant de processus de fond. Les données sont en accord avec les prédictions de fonds du modèle standard. Les résultats sont utilisés pour contraindre les paramètres de masse de plusieurs modèles de supersymétrie. Pour la majorité de ces modèles, les limites sont les plus strictes à ce jour.

La seconde recherche de nouvelle physique présentée dans la thèse considère, pour la première fois à CMS, des neutrinos lourds produits dans les désintégrations de mésons B. Cette analyse novatrice exploite un ensemble de données collectées en 2018 et contenant de l’ordre de 10 milliards d’événements $b\bar{b}$ pour une luminosité intégrée de 41.6 fb^{-1} . Les neutrinos lourds sont présumés avoir un long temps de vie afin de pouvoir étudier le cas où ils se mélangent faiblement avec les neutrinos du modèle standard. Aucun excès significatif de données vis-à-vis des fonds du modèle standard est observé. Des limites sont obtenues sous les hypothèses Majorana et Dirac pour différents paramètres décrivant la masse, le temps de vie, et le mélange des neutrinos lourds. Parmi ces contraintes, certaines sont les premières jamais présentées et d’autres sont les plus strictes à ce jour.

Enfin, la recherche de nouvelle physique ne serait pas possible sans d’excellents algorithmes pour reconstruire les particules. Le dernier projet présenté dans la thèse consiste en l’optimisation de l’un des algorithmes du calorimètre électromagnétique, réalisée en vue du Run 3 de prise de données. Toutes les analyses à CMS utilisant des données du Run 3 et comprenant des photons, électrons, jets, ou énergie manquante bénéficient de ce travail d’optimisation.

Acknowledgments

The accomplishment of the thesis work would not have been possible without the help, guidance, and support from a series of people. I would like to deeply and sincerely thank

Günther Dissertori, for allowing me to perform my PhD in such a great environment and to work on highly interesting research projects, for directing my PhD such that it reaches excellence, for encouraging me to accept opportunities and entrusting me with a leadership role with the BHNL analysis, for supporting me throughout the whole PhD and having always made time to listen and discuss, for believing in me since the beginning, and, above all, for his great humanity and unlimited kindness.

Jeff Richman, for his invaluable help with the convergence of the BHNL paper, for giving me so much of his time for discussing physics from a different perspective during our Thursday conversations, for helping me to take a step back and broaden my scope, and for his mentorship and willingness to support me as his own student.

Riccardo Manzoni, for his extraordinary guidance with the BHNL analysis, for being so generous with sharing his unlimited knowledge and his passion for physics, for his wit and humour, and all in all, for having become my CMS older brother.

Maria Giulia Ratti, for introducing me to the world of CERN and data analysis, for her supervision and guidance with the research projects, for having been always available, and for her care and attention to the advancement of the first half of my PhD.

Mauro Donegà, *Amina Zghiche*, and *Fabrice Couderc*, for their help on the clustering project and for having made of the clustering team a little ECAL family.

Slava Valouev, for his careful review of the BHNL analysis and for his support.

Mario Masciovecchio and *Bennett Marsh*, for their amazing work on the M_{T2} analysis.

Francesco Pandolfi and *Raffaella Tramontano*, for their contribution to the BHNL analysis, their availability, and their positivity.

Thea Arrestad, for her endless encouragements and for always sharing so generously her indestructible good energy.

My amazing IPA colleagues, especially *Kaustuv*, *Vasilije*, and *Jan*, for having been on board since the beginning of the PhD journey and for their friendship.

Daniel Treille and *Felicitas Pauss*, for accompanying me since my early steps as a physicist and for being a source of inspiration.

Finally, and most importantly, my ***friends*** and ***family***, for their love and resolute support.

Table of contents

Introduction	1
I Setting the scene	3
1 The success and limitations of the standard model	5
1.1 The standard model of particle physics	5
1.2 Neutrinos in the standard model	11
1.3 The physics at hadron colliders	12
1.4 Experimental tests of the standard model	14
1.5 The need for new physics	14
2 The experimental frame	17
2.1 The Large Hadron Collider	17
2.2 The Compact Muon Solenoid detector	19
2.3 Trigger and data acquisition system	23
2.3.1 The parking strategy	25
2.4 Particle reconstruction	26
2.5 Optimisation of the ECAL clustering algorithm	30
II Search for new physics in the supersymmetric sector	35
3 Main concepts of supersymmetry	37
3.1 The minimal supersymmetric standard model	37
3.2 Probing supersymmetry at hadron colliders	40
3.3 Simplified supersymmetric models	42
4 Search for supersymmetry using the M_{T2} variable	45
4.1 Search variables	47
4.1.1 The M_{T2} variable	47
4.1.2 Other kinematic variables	50
4.2 Event samples	50
4.2.1 Data samples	50
4.2.2 Simulated samples	52
4.3 Event reconstruction	52
4.4 Event selection	54
4.5 Event categories	56
4.6 Signal signatures	59
4.7 Background estimation	61
4.7.1 $Z \rightarrow \nu\bar{\nu}$ background	61
4.7.2 Lost-lepton background	68
4.7.3 QCD multijet background	73

4.8	Results and interpretation	73
4.9	Summary	82
III	Search for new physics in the neutrino sector	83
5	The physics of heavy neutrinos	85
5.1	Active neutrino oscillations	85
5.2	Heavy neutrinos as a solution to the neutrino mass puzzle	86
5.3	The neutrino minimal standard model	88
5.4	The phenomenological frame	89
5.5	Experimental landscape	91
6	Search for long-lived heavy neutrinos in the decays of B mesons	93
6.1	The B-parking data sample	96
6.2	Simulated signal samples	97
6.2.1	Heavy neutrino production and decay widths	98
6.2.2	Event generation	102
6.2.3	Signal signatures	104
6.3	Event reconstruction	106
6.4	Event selection and categories	108
6.4.1	Preselection	108
6.4.2	Categories	111
6.4.3	Selection	113
6.5	Monte Carlo corrections	119
6.6	Signal normalisation and parametrisation	122
6.6.1	Signal yield normalisation	123
6.6.2	Reweighting to arbitrary mixing scenarios	124
6.6.3	Reweighting to the Dirac-like scenario	125
6.6.4	Reweighting to arbitrary lifetimes	126
6.6.5	Measurement of the effective cross section for $B^\pm$ production	126
6.6.6	Signal shape parametrisation	128
6.7	Background estimation	129
6.7.1	Background sources	129
6.7.2	Standard model resonances	131
6.7.3	Background shape modelling	132
6.8	Systematic uncertainties	134
6.9	Results and interpretation	136
6.10	Summary	145
	Conclusion	147
	A Performance of the clustering optimisation	149
	Statement of originality	153
	List of acronyms	155
	List of figures	165
	List of tables	169
	Bibliography	181

Introduction

How can the mass of the neutrinos be so small? What causes the asymmetry in the abundance of matter and antimatter in the universe? What is dark matter made of? These fundamental questions remain unanswered in the standard model (SM) of particle physics, which is considered as the most complete and robust theoretical framework that describes the particles and their interactions. Possible extensions of the SM that provide solutions to these open questions are commonly referred to as *new physics*. The search for new physics is placed at the core of the research work conducted during the PhD thesis.

Two prominent classes of new physics models are investigated in the context of this work: supersymmetry (SUSY) and the neutrino minimal standard model (ν MSM). These models predict the existence of supersymmetric particles and of heavy neutrinos, respectively. The research projects carried out during the PhD aimed at discovering the presence of these hypothetical particles in the data collected by the CMS experiment at the CERN Large Hadron Collider (LHC).

In the first part of the thesis, the theoretical and experimental aspects needed to understand the research projects are described. In Chapter 1, an overview of the structure, achievements, and limitations of the SM is presented. In particular, we will highlight the important role played by the principle of symmetry in constructing robust theoretical foundations for the model. On the one hand, the preservation of the gauge symmetry provides a solid description of the strong, electromagnetic, and weak fundamental interactions. On the other hand, the process of spontaneous symmetry breaking is the building stone to explain how particles acquire a mass. These elements contribute to making of the SM one of the most successful theories of physics. However, we will discuss how the limitations of this model motivate the need for new physics.

The experimental apparatus, described in Chapter 2, used to conduct the research projects presented in this thesis is located at CERN, the largest laboratory for particle physics in the world. In particular, new physics is searched for in the data that are produced from proton-proton collisions obtained at a centre-of-mass energy of 13 TeV during the Run 2 of the LHC, and that are recorded by the CMS experiment. The data are analysed to identify the signatures of the processes involving the hypothetical particles that would allow them to be discovered.

The CMS detector is built with a series of subdetectors, each having a dedicated role in the detection and identification of the particles produced in the proton-proton collisions. The subdetector used to measure the energy of the photons and electrons is the electromagnetic calorimeter (ECAL). The first project presented in the thesis consists in optimising one of the ECAL reconstruction algorithms, in view of reaching stable and optimal performance during the Run 3 of data taking, characterised with an enhanced level of noise in the detector. The results from this work are summarised in Section 2.5.

The second part of the thesis revolves around the search for new physics in the supersymmetric sector. By postulating the existence of an additional symmetry between fermionic and bosonic states, SUSY extends the particle content with supersymmetric partners associated with each of the SM particles. This theoretical framework, described in Chapter 3, has been one of the most investigated new physics models, as it provides solutions to several of the limitations of the SM, among them the unknown nature of dark matter.

The search for new physics conducted during the thesis work in the supersymmetric sector is presented in Chapter 4. This search targets processes conserving the R -parity quantum number, defined in Section 3.1, in which the decay of pair-produced supersymmetric particles produced via the strong interaction give rise to all-hadronic final states with large transverse momentum imbalance. The analysis exploits a data set corresponding to an integrated luminosity of 137 fb^{-1} , which is the largest data set ever used so far at CMS for these purposes. The strategy of the search relies on the use of the M_{T2} variable, defined as a generalisation of the transverse mass for events in which pair-produced particles decay semi-invisibly, as it is sensitive to the different topologies between supersymmetric and SM events. An essential part of the work was the estimation of the two main sources of SM processes that give rise to similar signatures as the new physics processes.

Finally, the third and last part of the thesis is dedicated to the search for new physics in the neutrino sector. The fact that the SM neutrinos have a small, yet non-zero, mass remains, to date, unexplained and is one of the strongest evidences for the incompleteness of the SM. A possible mechanism for the generation of the neutrino masses involves new particles, referred to as heavy neutrinos. As summarised in Chapter 5, in the context of the ν MSM, these particles are not only a key to the neutrino mass puzzle, but they can also explain the asymmetry between matter and antimatter and be a candidate for dark matter.

The first search for long-lived heavy neutrinos in the decays of beauty mesons ever performed at CMS is presented in Chapter 6. Probing new physics particles in such decays is a breakthrough for this experiment, which is made possible by the collection, in 2018, of a special beauty-hadron-enriched data sample, referred to as the B-parking data sample, comprising of order $10^{10} \text{ b}\bar{\text{b}}$ events for an integrated luminosity of 41.6 fb^{-1} . This innovative search was fully initiated, designed, and conducted in the context of the thesis work. The analysis targets the signatures of heavy neutrinos that live long enough to travel a measurable distance inside the detector before decaying. While this requirement makes the analysis challenging, it maximises its sensitivity to, until then, unexplored regions of the parameter space. Moreover, as processes are considered in which the invariant mass of the heavy neutrino can be fully reconstructed, these hypothetical particles are probed with an unprecedented resolution in mass.

In summary, the present thesis reports on the three main projects carried out during the PhD: (i) the optimisation of one of the reconstruction algorithms of the ECAL in view of the Run 3 of data taking (Section 2.5), (ii) the search for supersymmetric particles using the M_{T2} variable (Chapter 4), and (iii) the search for long-lived heavy neutrinos in the decays of beauty mesons (Chapter 6). Throughout the document, each of the projects is motivated, described, and put into perspective within the current experimental landscape.

Part I

Setting the scene

Chapter 1

The success and limitations of the standard model

The standard model (SM) of particle physics is considered as the most solid theory that describes the fundamental constituents of matter and their interactions. Its success relies on that (i) it provides a unified framework explaining the strong, electromagnetic, and weak fundamental interactions, (ii) it is in excellent agreement with a wide range of experimental observations, and (iii) it has proved capable of predicting the existence of key particles that were subsequently discovered experimentally. A complete description of the SM is provided in Refs. [1, 2]. In this chapter, the theoretical concepts essential to understand the research work presented in the thesis are summarised. Section 1.1 gives an overview of the structure of the SM. Key elements to neutrino physics are then discussed in Section 1.2 and will serve as a baseline to motivate the extension of the SM in the neutrino sector, described in Chapter 5. The main aspects of hadron collider physics are then discussed in Section 1.3. The robustness of the SM is established through extensive experimental tests, a summary of which is given in Section 1.4. Nevertheless, despite this success, multiple limitations of the SM point towards the fact that it is not the ultimate theory for particle physics, and that it rather must be extended. These limitations, discussed in Section 1.5, motivate the need for new physics.

1.1 The standard model of particle physics

The standard model of particle physics is a renormalisable quantum field theory that describes the strong, electromagnetic, and weak interactions in a coherent framework with solid mathematical foundations. It is built on the *local gauge principle*, which implies that the Lagrangian describing the theory must be invariant under specific local transformations. This principle is at the core of the theory, as it dictates how the fundamental particles interact among themselves. In particular, it follows that the interactions between particles constituting matter, referred to as *fermions*, are themselves mediated by another type of particles, referred to as *gauge bosons*.

The gauge theory on which the SM is constructed is based on the $SU(3)_C \times SU(2)_L \times U(1)_Y$ symmetry group. The strong interaction, described by quantum chromodynamics (QCD), is based on a $SU(3)_C$ gauge symmetry that specifies the interactions of objects carrying so-called

colour charge. On the other hand, the electromagnetic and weak interactions are described within a common framework, known as the electroweak theory. This theory is based on the $SU(2)_L \times U(1)_Y$ symmetry, for which the weak isospin, I , and its third component, I_3 , as well as the weak hypercharge, Y , are conserved. The weak hypercharge is defined as $Y \equiv 2(Q - I_3)$, where Q is the electric charge, typically expressed in units of the elementary charge, e . While all the particles can interact weakly, only those carrying a non-zero colour or electric charge can interact strongly or electromagnetically, respectively.

Before summarising the main principles of QCD and of the electroweak theory and its spontaneous breaking, we first describe the particle content of the SM.

Particle content

The fundamental particles in the SM can be grouped into two classes of particles: the bosons and the fermions. They are defined based on the value of the *spin* of the particle, defined as its intrinsic angular momentum: bosons and fermions have integer and half-integer values of spin, respectively.

The gauge bosons, which have a spin equal to 1, are the quantum excitations of the gauge fields introduced to maintain the invariance of the Lagrangian density under the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge transformations. They are the mediators of the fundamental interactions:

1. The strong interaction is mediated by eight massless gluons, g .
2. The electromagnetic interaction is mediated by the massless photon, γ .
3. The weak interaction is mediated by the massive $W^\pm$ and Z^0 bosons, which are charged and neutral particles, respectively. In the following, these bosons are denoted with W and Z , respectively, without writing explicitly the electric charge.

The last known fundamental boson is the Higgs boson, which has, instead, a spin equal to 0. Unlike the aforementioned bosons, it does not arise directly from the local gauge principle, but rather from the spontaneous breaking of the electroweak symmetry, as discussed below.

At the time of writing, the SM comprises twelve fermions, consisting of six quarks and six leptons. For each fermion, there exists an antiparticle, which has the same properties as the particle, except for the charges that have opposite values. Because fermions have a spin equal to $1/2$, they can exist in two chiral states, referred to as *left-handed* and *right-handed* chiral states. All the fermions have been observed with the two chiralities, except for the neutrinos, which have been observed only with left-handed chirality, as discussed in Section 1.2.

The leptons and quarks can be both classified into three flavour families. For the quarks, the families are respectively composed of (i) the *up* (u) and *down* (d) quarks, (ii) the *charm* (c) and *strange* (s) quarks, and (iii) the *top* (t) and *beauty* (b) quarks. The u , c , and t quarks have an electric charge of $+2/3$ and are referred to as the *up-type* quarks, while the d , s , and b quarks have an electric charge of $-1/3$ and are referred to as the *down-type* quarks. Each quark carries a colour charge, while any bound state of quarks, referred to as a *hadron*, must be




Figure 1.1: Illustration of the particle content of the standard model.

a colour singlet. Hadrons formed with a triplet of quarks are denoted as *baryons*, while those formed from a quark-antiquark pair are denoted as *mesons*. The quarks are the only fermions sensitive to all three of the fundamental interactions, i.e. the strong, electromagnetic, and weak interactions.

The charged leptons do not carry colour charge and are therefore sensitive only to the electromagnetic and weak interactions. The three lepton families are based on the three flavours of the charged leptons: the electron (e), muon (μ), and tau lepton (τ). The corresponding neutral leptons, called neutrinos, are denoted as ν_e , ν_μ , and ν_τ . As further described in Section 1.2, they are sensitive only to the weak interaction, as they do not carry an electric charge.

The particle content of the SM is summarised in Fig. 1.1. In the following, the main ingredients on which the SM is based are further described.

Quantum chromodynamics

Quantum chromodynamics is the theory that describes the strong fundamental interaction. It is constructed based on the local gauge principle involving the invariance under $SU(3)_C$ local transformations, under which the quark field, ψ , must obey

$$\psi(x) \rightarrow \psi'(x) = e^{ig_s T_a \alpha_a(x)} \psi(x), \quad (1.1)$$

where g_s is the strong coupling strength, T_a ($a = 1, \dots, 8$) are the generators of the $SU(3)_C$ symmetry group, and α_a are functions of the space-time coordinate, x . The conserved charge associated with the invariance under such local transformations is the colour charge, C , which can take three values, commonly referred to as red, green, and blue. The only fermions that carry a non-zero colour charge are the quarks.




Figure 1.2: Measurements of the strong coupling constant, α_s , shown as a function of the energy scale, Q . The figure is taken from Ref. [3].

The invariance of the QCD Lagrangian under the local transformation in Eq. (1.1) is asserted with the introduction of eight gluon fields, $G_\mu^a(x)$. The associated particles, the gluons, are gauge bosons that carry both a colour and an anticolour charge, such that coloured fermions, i.e. quarks, can couple to them and hence interact strongly. Moreover, gluons are predicted to be massless, as the introduction of a gluon mass term to the Lagrangian breaks the local $SU(3)_C$ symmetry. The eight gluon fields, $G_\mu^a(x)$, correspond to the eight generators of the symmetry group. These generators do not commute, allowing for gluon self-interaction terms, giving rise to triple and quartic gluon vertices.

The existence of gluon self-interactions has important consequences. One of them is known as the *colour confinement*, which is a property unique to the strong interaction that explains how free quarks have never been observed experimentally. The colour confinement indeed implies that quarks cannot propagate as free particles, and are instead confined into colour singlet states, known as hadrons. As we will discuss in Section 1.3, the colour confinement has important implications for the experimental signatures of quarks and gluons.

The strong coupling constant, α_s , is defined as $\alpha_s \equiv g_s^2(4\pi)^{-1}$. Its magnitude is not constant and instead varies with the energy scale, Q , of the interaction, as shown in Fig. 1.2. At low energy scales, the value of α_s is large, and the quarks are strongly bounded into hadrons, while at high energy scales, the strength of α_s is much smaller, and the quarks can be treated as quasi-free particles. This latter regime is known as *asymptotic freedom*, and will be of relevance when discussing hadron-hadron collisions in Section 1.3.

Electroweak theory

The electromagnetic and weak interactions are described within a common model, known as the electroweak theory, which is based on the invariance of the Lagrangian under the $SU(2)_L \times U(1)_Y$ local gauge symmetry. The conserved charges associated with the symmetry are the weak isospin, I , and the weak hypercharge, Y . The charged-current and neutral-current weak interactions are mediated by the charged W bosons and by the neutral Z boson, respectively.

Because the W bosons are electrically charged, they couple together fermions whose electric charge differ by one unit. Therefore, it is convenient to organise the fermions into weak isospin doublets, $\eta(x)$, written as

$$\eta(x) = \begin{pmatrix} \nu_\ell(x) \\ \ell^-(x) \end{pmatrix}_L, \quad \ell = e, \mu, \tau, \quad (1.2)$$

for the leptons and as

$$\eta(x) = \begin{pmatrix} U(x) \\ D(x) \end{pmatrix}_L, \quad U = u, c, t \quad \text{and} \quad D = d, s, b, \quad (1.3)$$

for the quarks. The subscript L indicates the left-handed chirality for the doublets. Indeed, the weak interaction violates parity and thus does not treat equally the left- and right-handed chiral fields. In particular, the W bosons couple uniquely to left-handed particles and right-handed antiparticles. As a result, the right-handed particle and left-handed antiparticle chiral fields are not placed in weak isospin doublets but rather in weak isospin singlets with $I = 0$, such that they are unaffected by the $SU(2)_L$ local gauge transformations. These transformations are expressed as

$$\eta(x) \rightarrow \eta'(x) = e^{ig_W K_j \beta_j(x)} \eta(x), \quad (1.4)$$

where g_W is the weak coupling constant, K_j ($j = 1, \dots, 3$) are the three generators of the $SU(2)_L$ group, and $\beta_j(x)$ are three functions of the space-time coordinate, x .

Following from the local gauge principle, charged-current weak interactions are associated with the invariance under the local $SU(2)_L$ transformations in Eq. (1.4). In particular, the invariance of the Lagrangian is achieved with the introduction of three gauge fields, $W_\mu^{(1)}$, $W_\mu^{(2)}$, and $W_\mu^{(3)}$. The physical W bosons, associated with the gauge fields $W_\mu^\pm$, can then be obtained with the following linear combinations:

$$W_\mu^\pm = \frac{1}{\sqrt{2}} \left(W_\mu^{(1)} \mp W_\mu^{(2)} \right). \quad (1.5)$$

As opposed to charged-current weak interactions, the neutral-current weak interactions do not maximally violate parity, as the associated Z boson was observed experimentally to couple to both left- and right-handed particle chiral fields. Therefore, these interactions cannot arise from a theory based uniquely on the $SU(2)_L$ gauge group. Instead, they can be described within a theory that is invariant under an additional local symmetry, the $U(1)_Y$ gauge symmetry, which transforms a field $\psi(x)$ as

$$\psi(x) \rightarrow \psi'(x) = e^{ig' \frac{Y}{2} \gamma(x)} \psi(x), \quad (1.6)$$

where g' is the associated coupling constant and $\gamma(x)$ are functions of the space-time coordinate. Similarly to quantum electrodynamics, the theory describing the electromagnetic interaction for which the invariance under the $U(1)$ symmetry gives rise to a photon field that couples to the electric charge, Q , the invariance under the transformation in Eq. (1.6) gives rise to a gauge field, B_μ , that couples to the weak hypercharge, Y .

Within the electroweak theory, in which the Lagrangian is invariant under both the $SU(2)_L$ and $U(1)_Y$ symmetries, the photon and Z boson, described by the fields A_μ and Z_μ , respectively, can be expressed in terms of the neutral gauge fields, B_μ and $W_\mu^{(3)}$, as

$$\begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} B_\mu \\ W_\mu^{(3)} \end{pmatrix}, \quad (1.7)$$

where θ_W is the weak mixing angle. The electroweak unification implies that the couplings of the weak and electromagnetic interactions are related as

$$e = g_W \sin \theta_W = g' \cos \theta_W, \quad (1.8)$$

where e is the fundamental electric charge.

Similarly to the gluon fields in QCD, the addition to the Lagrangian of explicit mass terms for the electroweak gauge bosons, γ , W , and Z , breaks the invariance under the $SU(2)_L \times U(1)_Y$ symmetry. Therefore, the electroweak bosons are all predicted to be massless. While this is a consistent description for the photon, it is not for the weak gauge bosons, as they were found experimentally to be massive. A mechanism that allows for the weak gauge bosons to acquire a mass is thus needed for a consistent formulation of the standard model.

Electroweak symmetry breaking

Although QCD and the electroweak theory provide a solid description of the interactions between the fundamental particles, they cannot be the sole ingredients to build a complete and coherent theory describing particle physics. In fact, taken as is, these theories predict all the particles to be massless because (i) mass terms for the gauge bosons break the local gauge invariance of the theory, and (ii) mass terms for the fermions of the form $m\bar{\psi}\psi = m(\bar{\psi}_R\psi_L + \bar{\psi}_L\psi_R)$ are not allowed since the left- and right-handed chiral fields transform differently under the $SU(2)_L \times U(1)_Y$ symmetry. However, this prediction is in contradiction with the experimental observation that the weak bosons as well as the fermions (left aside for the moment the neutrinos, discussed in Section 1.2) are massive. Therefore, a mechanism must be derived such that it generates masses for the weak gauge bosons without breaking the local gauge principle, and explains masses of the fermions.

A solution was proposed by Englert, Brout, and Higgs [4, 5] and is commonly referred to as the *Higgs mechanism*. The minimal formulation of this mechanism relies on the introduction of a doublet of complex scalar fields, $\phi(x)$, chosen such that its potential presents an infinite number of degenerate minima, defined by

$$|\phi|^2 = \frac{v^2}{2}, \quad (1.9)$$

where v is the vacuum expectation value of the potential. The choice of the vacuum state, i.e. the state with the lowest energy, from the possible minima explicitly breaks the electroweak gauge symmetry. This process is referred to as *spontaneous symmetry breaking*, as the Lagrangian of the theory remains invariant, while the vacuum state does not. After the symmetry breaking, the field $\phi(x)$ is expanded about the vacuum state, giving rise to one massive scalar and three massless Goldstone bosons. The gauge is chosen such that the latter are absorbed into the physical weak bosons and that the field expansion can be written as

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \quad (1.10)$$

where $h(x)$ is the massive Higgs field that describes the Higgs boson. Through the Higgs mechanism, the weak gauge bosons acquire a mass equal to

$$m_W = \frac{1}{2}g_W v \quad \text{and} \quad m_Z = \frac{1}{2}v\sqrt{g_W^2 + g'^2}. \quad (1.11)$$

Therefore, the spontaneous breaking of the $SU(2)_L \times U(1)_Y$ symmetry allows for the generation of the mass of the weak gauge bosons without breaking the local gauge invariance of the model.

The fermions can also acquire a mass when they couple to the Higgs boson. The coupling, λ_f , of a fermion f to the Higgs field, called Yukawa coupling, allows for the introduction of mass terms of the form of

$$\mathcal{L}_{\text{mass}} = -\frac{1}{\sqrt{2}}\lambda_f(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L)(v + h). \quad (1.12)$$

The mass, m_f , of the fermion f is equal to

$$m_f = \frac{1}{\sqrt{2}}\lambda_f v. \quad (1.13)$$

While the charged leptons and the quarks become massive through their coupling to the Higgs field, it is not the case for the neutrinos, as discussed in the next section.

1.2 Neutrinos in the standard model

Neutrinos in the standard model, also referred to as *active* neutrinos, are fermions that carry neither strong nor electromagnetic charges, and therefore only interact weakly. The conservation of lepton number in all observed neutrino processes other than neutrino oscillations shows that the three neutrino states involved in the charged-current weak interactions are distinct [1]. As a result, three neutrino states, ν_e , ν_μ , and ν_τ , are introduced, such that they correspond to the weak eigenstates produced in charged-current weak interactions along with the leptons e , μ , and τ , respectively. The fact that there exist three different neutrino states is confirmed by the studies on the Z production in e^+e^- collisions at the Large Electron Positron collider (LEP) at CERN, which measured that the number of active neutrinos was equal to $N_\nu = 2.9963 \pm 0.0074$ [3].

The neutrino weak eigenstates, ν_ℓ , are defined as the $SU(2)_L$ partners of the charged lepton mass eigenstates, ℓ , and are therefore part of the lepton doublets defined in Eq. (1.2). Unlike the other fermions, the active neutrinos have been experimentally observed, so far, only with

left-handed chirality [6]. In the minimal SM, the non-observation of right-handed chiral neutrino states is interpreted as the non-existence of right-handed neutrinos. A direct, and major, consequence is that the active neutrinos cannot acquire a mass via the same process as for the other fermions, i.e. via the Yukawa coupling to the Higgs doublet, as can be inferred from Eq. (1.12). Furthermore, as argued in Ref. [7], bare mass terms of the form $\bar{\nu}_L \nu_L^c$ cannot arise from loop corrections as they do not conserve the total lepton number symmetry. Therefore, in the SM, neutrinos are considered to be massless particles. We will discuss in Section 1.5 that this prediction is contradicted experimentally, motivating the need for new physics.

1.3 The physics at hadron colliders

The searches for new physics carried out in the thesis work were conducted at CMS, one of the experiments at the CERN LHC, which is a hadron collider. In this section, the main aspects of hadron collider physics are presented, with a focus on proton-proton collisions.

The physics at a hadron collider is determined by the fact that the particles that are accelerated and collided (i) are heavier than electrons; (ii) are composite, i.e. have an internal structure rather than being point-like; and (iii) have coloured constituents. On the one hand, this makes of hadron colliders a favourable environment for discoveries, as (i) the lower synchrotron radiation compared to lepton colliders makes it possible to reach higher centre-of-mass energies, (ii) the available energy for the interactions spans a wide range of values rather than being fixed, and (iii) the new physics particles can be produced via the strong interaction. On the other hand, hadron collisions give rise to complex processes and final states resulting from the convoluted structure of the protons. As a consequence, the associated theoretical computations, object reconstruction, and data analysis can reach high levels of difficulty.

Protons are indeed not elementary particles and have instead an internal structure, composed of quarks and gluons, referred to as *partons*. In particular, protons are composed of two up and one down valence quarks, as well as a sea of gluons and quark-antiquark pairs. The underlying structure of the proton is described with the *parton distribution functions* (PDF), which give the probability for a given parton to carry a fraction x of the proton momentum.

The physics governing proton-proton collisions is described by QCD. When the two hadrons collide, their partons can be considered as quasi-free, as the high energy scale reached by the collision results in the regime of the asymptotic freedom, defined in Section 1.1. As a consequence, proton-proton collisions are actually parton-parton collisions that have the following properties:

- The two partons carry a fraction x_1 and x_2 of the momentum of the respective protons. Because these fractions are unknown, the actual momentum of the colliding parton-parton system cannot be resolved in the direction of the beamlines. On the other hand, the conservation principle implies that the momentum of the initial and final states projected onto the plane transverse to the direction of the colliding protons are both null. For that reason, the relevant physics at hadron colliders takes place in the transverse plane.
- The parton interactions involving a large momentum transfer are referred to as the *hard*




Figure 1.3: Illustration of a hard scattering process in a hadron-hadron collision.

scattering. Such processes typically carry interesting physics content, as the particles produced via the deep inelastic scattering of the protons typically have large transverse momenta.

- Besides the hard parton interaction, soft processes may also take place. For instance, partons involved in the hard process may emit gluons, giving rise to initial or final state radiations. Moreover, the partons that do not give rise to the hard scattering may nevertheless collide, resulting in soft scattering processes constituting the so-called *underlying event*.

An example of a hard scattering process, $h_1 h_2 \rightarrow AB$, is illustrated in Fig. 1.3. The cross section of this process for a typical energy scale, Q , can be computed as the convolution of the parton-level cross section, $\hat{\sigma}^{ij \rightarrow AB}$, with the PDF, f , of the hadrons:

$$\sigma^{h_1 h_2 \rightarrow AB} = \int_0^1 dx_1 \int_0^1 dx_2 \sum_{i,j} f_i(x_1, \mu_F^2) f_j(x_2, \mu_F^2) \hat{\sigma}^{ij \rightarrow AB}(Q^2, \mu_R^2, \mu_F^2). \quad (1.14)$$

Two energy scales are introduced: the renormalisation scale, μ_R , regulates the possible ultra-violet divergences arising from the loop corrections to the cross section, while the factorisation scale, μ_F , controls the infrared and collinear divergences.

Because of the colour confinement, introduced in Section 1.1, quarks and gluons cannot be observed as individual particles. Therefore, the experimental signature of a parton is a cluster of particles collimated along the direction of flight, called a *jet*. Indeed, as a parton is produced at high energy, it can give rise to a parton shower, resulting from the successive radiation of quarks to gluons, and of gluons to quark-antiquark pairs. When the energy is sufficiently low that no further radiation process is possible, the existing partons combine into hadrons. This process is called *hadronisation*. Jets are hence defined as the collection of the particles produced during the hadronisation process of a parton. The algorithm used to reconstruct jets in the context of this work is discussed in Section 2.4.

1.4 Experimental tests of the standard model

Quantum chromodynamics, the electroweak theory, and the Higgs mechanism are the key ingredients that constitute the current formulation of the standard model of particle physics. A model based on these elements has indeed proven to be consistent with the experimental observations up to very high precision.

Experimental tests play a key role for establishing a theory. Not only do they serve to verify that the theory describes the data correctly, but they also allow the free parameters of the theory to be constrained. In the minimal formulation of the SM, there are 19 such free parameters, corresponding to (i) the Yukawa couplings, (ii) the coupling constants, (iii) the vacuum expectation value of the Higgs potential and the Higgs mass, and (iv) the parameters describing the mixing of quarks.

Over the past decades, a tremendous amount of tests in the strong, electroweak, and Higgs sectors have been carried out in many experimental facilities. An overview of the results from these tests is presented in Ref. [3] and references therein. The SM has shown an extraordinary capability to describe the observed data to a very high level of accuracy and for an impressively large range of processes. This achievement makes of the SM one of the greatest successes of modern physics.

A striking example illustrating the success of the SM is shown in Fig. 1.4. The predicted SM cross sections for a large number of processes are found to be in excellent agreement with the measured values from the CMS experiment. The fact that various processes are considered, including inelastic proton-proton collisions and the production of jets, top quarks, and Higgs bosons, and that fourteen orders of magnitude in the value of the cross section are covered, shows that the SM is capable of explaining a wide range of physics processes.

The strength of the SM does not only rely on that it can describe the observed data well, but also on that it is a highly predictive theory. For example, weak neutral-current interactions, mediated by the Z boson, were predicted within the electroweak theory before they had been observed experimentally. The discovery of the Z boson at the UA1 [9] and UA2 [10] experiments at the CERN Super Proton Synchrotron constituted a solid validation of the electroweak theory. Another successful prediction of the SM is the existence of the Higgs boson. The discovery of a particle compatible with the Higgs properties in 2012 by the CMS [11] and ATLAS [12] experiments at the CERN LHC is considered as a major accomplishment of the SM.

The SM is unarguably the most solid and well-tested theory for particle physics. However, it is not considered to be a complete theory, as it does not provide explanations for fundamental questions, as discussed in the next section.

1.5 The need for new physics

While the standard model provides a solid theoretical framework describing the strong, electromagnetic, and weak interactions between fundamental particles, it does not incorporate a description of gravity. This fact is often considered to be the main indication that the SM is a




Figure 1.4: Comparison between the predicted value for the cross section of various SM processes with the value obtained from measurements performed at the CMS experiment at the CERN LHC. The figure is taken from Ref. [8].

low-energy effective theory, and that it must therefore be extended. The idea that the SM is not a complete theory is reinforced by that it does not explain a series of phenomena:

- The hierarchy problem. The computation of the mass, m_{H} , of the Higgs boson must include corrections from fermion loops to the bare mass [13]. These corrections, δm_{H} , are of the form

$$\delta m_{\text{H}} = \frac{|\lambda_f|^2}{8\pi^2} \Lambda^2, \quad (1.15)$$

where λ_f is the Yukawa coupling of the fermion and Λ is the scale up to which the SM is valid. This scale is typically chosen to be the Planck scale, corresponding to an energy of $\mathcal{O}(10^{19})$ GeV. The mass of the Higgs boson being measured to be about 125 GeV therefore lies on a very different energy scale than Λ . For the calculation to agree with the measurement, the bare mass of the Higgs boson must thus be significantly fine-tuned. This fine-tuning makes the theory unnatural and hints to the existence of a mechanism preventing it. Such mechanism is provided in a supersymmetric theory, as discussed in Section 3.1.

- **Neutrino masses.** As discussed in Section 1.2, the neutrinos are considered to be massless particles in the SM. However, the observation of the neutrino oscillation phenomenon implies that at least two neutrinos are massive [14]. A mechanism for the mass generation of the neutrinos is currently missing in the SM and could follow from the addition of sterile right-handed neutrinos to the model, as described in Section 5.2.

- Asymmetry in the abundance of matter and antimatter. While all the particles and antiparticles should have annihilated after the Big Bang, during which they were produced with the same abundance, the universe is however clearly dominated with particles [15]. A mechanism that explains how matter won over antimatter is currently lacking in the SM. As discussed in Section 5.3, the oscillations between heavy neutrinos provide a possible answer to this puzzle.
- Dark matter. Several astronomical and cosmological observations suggest the existence of dark matter, which is constituted with one or several particles other than these described within the SM [16]. The nature of dark matter remains unknown, and could consist for instance of supersymmetric particles or heavy neutrinos.

The fact that these questions are left unanswered in the SM indicates that it is an incomplete theory and that it must be extended with new physics models. While many of these models have been theoretically derived, they must be tested experimentally to be either confirmed or rejected. There are two main strategies to search for the manifestation of new physics in the LHC data:

1. Indirect searches: very precise measurements of SM processes are conducted. Deviations from the expectation are interpreted as arising from contributions from new physics processes.
2. Direct searches: the presence of a new particle predicted in a new physics model is searched for based on the expected signatures from its decay. The process that is investigated is referred to as *signal*, while SM processes that mimic the signal signatures are referred to as *background*. An excess of data over the background prediction would be a sign for the presence of a new particle in the data.

In the work performed for this thesis, the latter approach is employed to search for new physics. Two different new physics models are considered, either predicting the existence of supersymmetric particles or of heavy neutrinos. The presence of these particles is searched for in the data collected by the CMS experiment at the LHC, described in the next chapter.

Chapter 2

The experimental frame

The new physics searches conducted in this work are performed using the infrastructure at CERN, the European Organisation for Nuclear Research, which is the largest particle physics laboratory in the world. The experimental apparatus consists of a particle accelerator, the Large Hadron Collider (LHC), as well as a particle detector, the Compact Muon Solenoid (CMS). In this chapter, an overview of the main experimental design aspects is presented. Section 2.1 describes the LHC, while Sections 2.2 and 2.3 provide an overview of the CMS detector and of its trigger system, with a particular focus on the data parking strategy. The main elements to reconstruct the particles entering the final states of the new physics searches are then summarised in Section 2.4. Finally, an important project carried out during the thesis is the optimisation of one of the reconstruction algorithms using the signals from the electromagnetic calorimeter. The main results from this optimisation are presented in Section 2.5.

2.1 The Large Hadron Collider

The LHC is currently the largest and highest-energy particle collider ever constructed. Placed in the former tunnel of the Large Electron Positron (LEP) collider, this synchrotron has a circumference of about 27 km. The machine primarily collides protons with protons, although collisions of heavy ions are also carried out. Because the thesis work was performed using proton-proton collisions, only this mode is considered in the following. A detailed description of the LHC is presented in Ref. [17].

The LHC was built with the aim of giving the possibility to cover an unprecedentedly large physics program. The main physics goals are (i) in-depth studies of SM processes; (ii) studies of very rare processes; (iii) the search for particles predicted in the SM, such as the Higgs boson; (iv) and finally the search for new physics particles, which is the main goal of the thesis work.

The particle accelerator was designed to collide particles at a centre-of-mass energy, $\sqrt{s}$, of up to 14 TeV, which is a factor of about 7 higher compared to the previous record of energy obtained with the Tevatron at Fermilab [18]. Colliding particles at high energy is required for the production of heavy particles. Moreover, the LHC was designed such that it would reach a peak instantaneous luminosity, L , of order $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. The high luminosity is a key component to

Table 2.1: List of runs of data taking that have taken place up to the moment of writing. For each run, the time lapse, the maximal centre-of-mass energy, as well as the delivered integrated luminosity are reported. The asterisk (*) symbol indicates that the values are reported as of the time of writing: the Run 3 of data taking is expected to last until 2026, and reach an integrated luminosity of about 160 fb^{-1} .

	Years	Maximal $\sqrt{s}$ (TeV)	Delivered $\mathcal{L}$ (fb^{-1})
Run 1	2009 - 2013	7	30
Run 2	2015 - 2018	13	164
Run 3	2022 - 2024*	13.6	124*

make the study of very rare processes possible. Indeed, for a given process with cross section σ , the number, N , of produced *events* is obtained as

$$N = \sigma \int L dt = \sigma \mathcal{L}, \quad (2.1)$$

where the time-integrated luminosity, $\mathcal{L}$, was introduced. An event is defined as the result of the fundamental interaction between the partons of the colliding protons.

The LHC started operating in 2009 and is expected to be in service until the 2040s. The data taking is organised in a series of *runs*, separated by non-running periods, referred to as the *long shutdown* periods, during which the accelerator as well as the detectors are maintained and upgraded. The period of each run that has taken place up to the time of writing is given in Table 2.1 together with the corresponding maximal centre-of-mass energy and delivered integrated luminosity. The data analysed in the thesis work to search for new physics were collected during the Run 2 of data taking, while the optimisation work on the particle reconstruction algorithm, presented in Section 2.5, was performed in view of the Run 3 of data taking. After the long shutdown following the end of Run 3, during which the accelerator will go through a major upgrade, the LHC will enter the High-Luminosity phase, during which the total integrated luminosity is expected to be increased by a factor of about 10 [19].

After the protons are obtained from the ionisation of hydrogen atoms, they are accelerated in different stages before being injected in the LHC. The chain of accelerators in the CERN complex is illustrated in Fig. 2.1. The protons are successively accelerated by (i) a linear accelerator, (ii) the proton synchrotron booster, (iii) the proton synchrotron, and (iv) the super proton synchrotron, which make them reach energies of 0.05, 1.4, 25, and 450 GeV, respectively. The protons are then passed from the super proton synchrotron to the LHC in so-called *fills* that last around 8 hours. During a fill, the protons are injected in two separate beams in the form of bunches containing of order 10^{11} protons and spaced from one another by 25 ns. The two proton beams contain about 2800 proton bunches and collide with a frequency of about 40 MHz. The beams are constrained to a circular path and are focused by superconducting dipole and quadrupole magnets operating at a temperature of 1.9 K and delivering a magnetic field up to 8.3 T.

The two proton beams circulate in opposite directions in parallel beam pipes, and collide at four points around the LHC ring, at which detectors are placed to record and study the activity




Figure 2.1: Illustration of the CERN accelerator complex. The figure is taken from Ref. [24].

produced by the collision. The four main LHC detectors are ATLAS [20], CMS [21], LHCb [22], and ALICE [23]. The thesis work was conducted at the CMS experiment, which is described in the following.

2.2 The Compact Muon Solenoid detector

The CMS detector is designed to detect the products of the LHC collisions and to measure the properties of the underlying processes at high precision, in view of studying various physics phenomena. It is one of the two general-purpose experiments at the LHC, the other being ATLAS, and it is capable of investigating a large range of physics processes. One of the main achievements attained at the CMS experiment is the discovery of the Higgs boson in 2012 [11].

The coordinate system employed by the CMS experiment is shown in Fig. 2.2, where the spatial coordinates, x , y , and z are defined together with the azimuthal, ϕ , and polar, θ , angles. The position of the objects is often identified based on ϕ and η , where the pseudorapidity η is defined as $\eta \equiv \ln \tan(\theta/2)$. Moreover, the longitudinal direction is defined as the direction along the z -axis, while the transverse direction is defined as the $x - y$ plane.

An illustration of the CMS detector is shown in Fig. 2.3. With a size of 21.6 m in length and 14.6 m in diameter, it is made of one central section, the *barrel*, surrounding the beam axis, and




Figure 2.2: Coordinate system used at the CMS experiment.

of two disks, the *endcaps*, perpendicular to the beam axis and closing the barrel at both ends. Among the main characteristics of the detector are (i) a superconducting solenoid providing a strong magnetic field of 3.8 T, (ii), a high-quality tracking system, (iii) a high-resolution electromagnetic calorimeter, and (iv) a robust muon spectrometer. A detailed description of the CMS detector can be found in Ref. [21].

While the LHC offers the unique opportunity to study physics produced at a high centre-of-mass energy and at a high luminosity, it also creates a challenging environment for the detector to operate. To face these challenges, an impressive amount of technology was gathered to build the detector. For instance, the different detector parts were designed with a high granular-




Figure 2.3: Illustration representing the CMS detector. Human figures are drawn to highlight the size of the detector. The figure is taken from Ref. [21].




Figure 2.4: Illustration of a slice of the CMS detector. The interaction of several SM particles with the different subdetectors is represented by the different coloured lines. The figure is taken from Ref. [25].

ity, such that the detector can function with a large number of simultaneous interactions per bunch crossing, referred to as *pileup*, resulting from the high beam crossing rate and intensity. Moreover, the detector elements were developed such that they can operate in a high radiation environment. Finally, an efficient trigger system was designed to cope with the large event rate following from the high luminosity.

The CMS detector is composed of a series of subdetectors, placed concentrically around the interaction point. These subdetectors are (i) the tracking system, (ii) the electromagnetic calorimeter (ECAL), (iii) the hadron calorimeter (HCAL), and (iv) the muon spectrometer. Each of them is dedicated to the detection and reconstruction of particles of different nature. An illustration of a slice of the CMS detector is given in Fig. 2.4, together with a representation of the interaction of several SM particles with the different subdetectors, further described in Section 2.4. The role of each subdetector and of the magnet is summarised in the following, where references presenting complete descriptions of the different elements are provided. The different constituents are presented in order of increasing distance with respect to the interaction point.

Tracking system

The tracking system [26] aims to reconstruct (i) the trajectory of charged particles, (ii) the primary vertex of the parton collision, and (iii) the secondary vertices from the decay of long-

lived particles or from pileup collisions. Because this system operates in the strong magnetic field produced by the superconducting solenoid, the momentum of the charged particles can be reconstructed with high resolution from their trajectories.

This system, covering the pseudorapidity region $|\eta| < 2.5$, is made of two parts: (i) the pixel detector, composed of silicon sensors organised in 4 layers in the barrel and 3 disks in the endcaps; and (ii) the strip detector, comprising silicon microstrips organised in 10 layers in the barrel and 11 disks in the endcaps. The two detector parts provide excellent spatial resolution, reaching values in the radial direction as low as 10 and 20 μm , respectively. Moreover, the silicon is a material with strong resistance to radiation, which is an important requirement for a robust tracking system, given that it is placed directly around the beam pipe. Additionally, the tracking system was designed with a minimal amount of material, such as to reduce the possible interactions with the tracker of the particles before they reach the calorimeters.

The pixel detector was upgraded in 2016, during which the number of channels was almost doubled [27], improving the resolution of the secondary vertices reconstruction and of the impact parameter measurement. A high resolution is indeed needed for the purposes of the identification of jets originating from b quarks, such as those used in the search described in Chapter 4, and the identification of long-lived particles, placed at the core of the search described in Chapter 6.

Electromagnetic calorimeter

The ECAL [28] is dedicated to the precise measurement of the energy of the electrons and photons. The energy is measured through the collection of the scintillating light from electromagnetic showers arising from the succession of Bremsstrahlung and pair production processes induced by the interaction of these particles with the ECAL material.

The ECAL is a homogeneous calorimeter consisting of scintillating crystals made of lead tungstate (PbWO_4) and of photomultipliers used to amplify the scintillation light. The choice of PbWO_4 as the crystal material was motivated notably by (i) its high density (8.28 g/cm^3); (ii) its short radiation length (0.89 cm), controlling the longitudinal spread of the shower; and (iii) its small Molière radius (2.2 cm), controlling the transverse spread of the shower. As a result, the shower originating from one particle is almost entirely contained within a few crystals. Moreover, its short scintillation time implies that 80% of the light is collected before the crossing of the next LHC bunch. The ECAL is a calorimeter with fine granularity that provides an excellent energy resolution, down to 1% for electrons or photons with energy above 10 GeV. Finally, another sampling calorimeter, the pre-shower, is placed in front of the ECAL at the endcaps, to improve the identification of neutral pions decaying to two photons.

The optimisation of one of the ECAL reconstruction algorithms was one of the projects carried out during the PhD thesis. It is described in Section 2.5.

Hadron calorimeter

The HCAL [29] is the CMS subdetector in charge of absorbing hadrons and of measuring their energy. This measurement is particularly needed for the reconstruction of jets and missing trans-

verse momentum, p_T^{miss} , defined in Section 2.4.3. When hadrons interact with the HCAL, they produce a shower resulting from electromagnetic and strong processes. The shower is collected within a succession of alternating layers of brass absorber and plastic scintillator constituting the sampling calorimeter. While most of the HCAL constituents are placed between the ECAL and the magnet, a small outer calorimeter is placed outside of the solenoid, to allow for the collection of the entire shower produced by very-high energetic particles. Finally, two additional hadron calorimeters, referred to as the very-forward calorimeters, are placed in the forward region of CMS and are useful for the reconstruction of particles produced quasi-collinearly with the beam direction.

Superconducting magnet

A superconducting solenoidal magnet is placed at the core of the CMS detector. Operating at a temperature of 4.5 K, it is able to generate a uniform magnetic field of 3.8 T within its volume. The strong magnetic field bends the trajectory of the charged particles and thus allows their transverse momenta to be measured with high precision. The magnetic flux is returned through an iron yoke placed outside of the magnet. The elements of the muon spectrometer are embedded within the yoke.

Muon spectrometer

Because of the high penetration power of muons, which prevents them from being stopped by the tracking system, the calorimeters, and the magnet, a dedicated system to measure their properties is needed. This system, referred to as the muon spectrometer [30], is placed outside of the superconducting magnet, within the iron return yoke providing an average magnetic field of 1.8 T and allowing for a second measurement of the bending radius for muons.

The muon spectrometer comprises three types of gaseous detectors, detecting the presence of a particle based on gas ionisation. Drift tubes (DTs) are placed in the region $|\eta| < 1.2$, while cathode strip chambers (CSCs) are placed in the region $0.9 < |\eta| < 2.4$. Both the DTs and the CSCs provide an excellent spatial resolution, as low as $40\ \mu\text{m}$, which allows for a precise measurement of the muon transverse momentum. Additionally, the robustness of the system is reinforced by resistive plate chambers (RPCs), placed in the region $|\eta| < 1.6$. Because RPCs have an excellent time resolution, they are suitable for the purposes of event triggering, discussed in the next section.

2.3 Trigger and data acquisition system

As anticipated in Section 2.1, the LHC makes it possible to collide protons at a rate of about 40 MHz. As the collection of the readout from all the detector channels corresponds to a size of about 1 MB, it is technically not feasible to store each event produced in the collisions. A key component of the CMS experiment is therefore the design of the *trigger system*, which selects the events to be stored and subsequently analysed. As the events that are not selected

are lost, it is of primary importance that the trigger conditions are optimised such that events interesting in view of the CMS physics program are kept. On the other hand, events should be discarded if they arise from soft collisions of partons, as they constitute an overwhelming source of background. A major challenge in the design of the trigger system is hence to select events as inclusively as possible, while managing the data rate to controllable levels.

The CMS trigger system has a two-level structure comprising

1. The Level 1 (L1) trigger [31].
2. The High Level Trigger (HLT) [32].

The L1 trigger is hardware-based, while the HLT is software-based. They have in common that they select events based on their topology and particle content, although they do not exploit the same amount of information about the event. Indeed, the L1 trigger must process more events than the HLT, within a shorter latency, which implies that it can rely only on a reduced content of the reconstructed event. Instead, more sophisticated algorithms can be used at the HLT. Overall, the trigger system reduces the event rate from 40 MHz to about 1 kHz.

The L1 trigger is composed of custom hardware processors placed on the detector. As the LHC collides proton bunches every 25 ns, the L1 trigger must be designed such that it is able to make the decision on whether to keep an event or not within a very short time period. The decision process is segmented into different logical rules that can be run in synchronous pipelines, such that the latency time is fixed to about 4 μ s. Because of this short time lapse, not all the information collected by the front-end electronics can be used in the decision process. Instead, the properties of the particles in an event are reconstructed with algorithms using coarse data from the calorimeters and the muon detectors only. An event is selected in case one or several of the associated particles meet predefined criteria. Furthermore, the overall event rate selected by the L1 trigger must be of the order of 100 kHz. In case too many events meet the selection criteria, a so-called *prescale* can be applied, such that only a fraction of these events is ultimately selected.

After the events are accepted by the L1 trigger, they are processed through the HLT, which further reduces the event rate down to a few kHz. The HLT trigger is based on software and is implemented on a computer farm. Because fewer events are processed by the HLT than by the L1 trigger, the time allocated to make the trigger decision at the HLT is significantly larger, corresponding to approximately 200 ms per event. Therefore, more complex algorithms than at the first trigger level can be used to select events. Indeed, particles are reconstructed using reconstruction algorithms that are similar to those used for the offline data analysis, optimised for fast processing. Therefore, the full detector readout, including the information from the tracking system, is exploited. The events are selected if they contain one or several particles passing given kinematic thresholds, the list of which is referred to as an HLT *path*. Typically, the more stringent are the thresholds, the lower is the rate of recorded events. In the context of this work, the data analysed for the search for supersymmetric particles were collected with HLT paths imposing requirements on jets and/or missing transverse momentum, while these

analysed for the search for heavy neutrinos were collected with HLT paths imposing thresholds on muons, as discussed in the respective chapters.

After the events are selected by the trigger system, they are stored and subsequently reconstructed using the full offline CMS reconstruction software, which makes use of the most advanced, and often computationally expensive, algorithms to obtain the most refined objects for data analysis. In the normal data stream, the events are promptly reconstructed, i.e. they are processed through the reconstruction software up to about 48 hours after having been recorded. However, an alternative strategy, known as the *parking* strategy, makes use of different data flow such that CMS can cope with an increased event rate. This strategy is discussed in the following.

2.3.1 The parking strategy

As a clear sign for new physics has not been discovered yet, the need for the design of novel strategies to provide sensitivity to new types of signals in the data becomes more and more pressing. The implementation of such strategies can in some cases take place at the level of the trigger. Indeed, as mentioned above, events that are not selected by the CMS trigger system are not recorded and are therefore lost. Therefore, alternative trigger strategies are interesting in that they can give access to different regions of the kinematic phase space, in which new physics signals may be comprised.

Various new physics models predict the existence of new particles with low masses and feeble couplings to the SM [33]. Such particles have decay products with low momenta and may thus be discarded by the trigger system. Lowering the values of the trigger thresholds on the particle properties would however make the recording of these events possible. Yet, lowering the value of the thresholds unavoidably increases the event rate. As a consequence, specialised data processing techniques must be designed to handle these larger rates.

One of these techniques is known as the data parking strategy [34]. It consists in delaying the reconstruction of the events to the point in time in which the computational resources become available, instead of having the events to be reconstructed promptly. With this strategy, the limitation imposed by the finite resources to reconstruct events is overcome, and the collection of large event samples recorded using trigger paths with low thresholds is made possible. As a consequence, the sensitivity of CMS in rare phenomena and low-mass new physics particles is enhanced.

The *B-parking* data set [34, 35] is an example that shows the success of the CMS parking strategy. Collected with triggers targeting muons originating from the semileptonic decays of B hadrons, this sample offers of the order 10^{10} $b\bar{b}$ events with a purity of about 80%. The events were accumulated at the end of Run 2 in 2018 and were reconstructed using the parking strategy in 2019 during the long shutdown. This data set gives the opportunity for CMS to conduct precise and competitive studies in B physics, although the detector was primarily not designed for such purposes. While this sample was originally collected with the aim to perform precise tests of the lepton flavour universality [36], its scope was extended beyond this objective.

Indeed, the B-parking sample made possible the search for long-lived heavy neutrinos in the decay of B mesons, designed during the thesis work and presented in Chapter 6. This work is the first search for new physics performed using this special data sample and has motivated the initiation of other searches by proving that world-best results could be obtained with these data. Further information on the trigger composition is presented in Section 6.1.

In light of the positive achievements obtained during Run 2 with the B-parking data set, this program was extended during the Run 3 of data taking [34]. In particular, two sets of trigger algorithms, either selecting dimuon or dielectron events, were developed to record B hadron decays. Furthermore, the parking program was diversified to cover a larger range of physics processes. Among them, triggers targeting the pair-production of Higgs bosons and targeting the decay of long-lived particles exploit the parking strategy.

2.4 Particle reconstruction

After an event is selected by the trigger system, its content must be fully reconstructed to measure the properties of the particles, such that the physics of the underlying processes can be studied. When a particle passes through the detector, it may leave signals in the various subdetectors, as illustrated in Fig. 2.4. Typically, signals in the tracking and muon systems are referred to as *hits*, while signals in the ECAL and HCAL are referred to as *energy clusters*. The collection of these signals allows for the reconstruction of the particle properties.

As explained in Section 1.3, the transverse momentum, $\vec{p}_T$, of the particle, defined as the projection of the full momentum, $\vec{p}$, on the transverse plane, is among the useful properties that are measured. In the following, the magnitude of the vector $\vec{p}_T$ is denoted as p_T . Other useful properties of the particles are for instance the coordinates in the $\eta - \phi$ plane and the longitudinal and transverse impact parameters, $|d_z|$ and $|d_{xy}|$, of the associated tracks.

Placed at the core of the CMS event reconstruction, the particle-flow (PF) algorithm [25] aims to reconstruct and identify each object in an event by combining the information collected in the different subdetectors. First, muons are reconstructed from the hits in the silicon tracker and in the muon spectrometer. Then, the energy clusters in the ECAL are used to reconstruct electrons and photons. While the latter are neutral particles and do not leave a track in the tracking system, the former are reconstructed from the combination of a track with the ECAL energy clusters. In a similar way, charged (neutral) hadrons are reconstructed based on the energy deposits in the ECAL and HCAL, with (without) the combination with the information from the tracking system. Finally, the combination of the identified particles allows for the reconstruction of jets and missing transverse momentum, $\vec{p}_T^{\text{miss}}$, as further discussed below.

The objects present the final state of the searches conducted in the thesis work are given the main focus. These objects are the jets and missing transverse momentum for the search for supersymmetry described in Chapter 4, and the muons, electrons, and pions in the search for heavy neutrinos described in Chapter 6. Photons are also briefly described as they are at the core of the optimisation work described in Section 2.5.

2.4.1 Tracks and vertices

Charged particles traversing the tracking system create signals in the sensors along their trajectory. These signals can be clustered into hits, which are the building blocks to reconstruct the track trajectory using the combinatorial track finder algorithm [37], which relies on Kalman filter techniques [38].

The final state in the search for heavy neutrinos presented in Chapter 6 comprises a charged pion. It is identified as a reconstructed track with high purity. Moreover, the mass of the pion is used when computing the four-momentum of the object.

The vertices are defined as the location of the proton-proton interaction points in the event. The primary vertex (PV) corresponds to the vertex from the hard interaction. Other vertices in the event can either arise from the decay of long-lived particles, such as B mesons, or from pileup collisions.

The vertices are reconstructed by aggregating the reconstructed tracks in the event using a dedicated algorithm [39]. The tracks are clustered based on their coordinate on the z -axis at their point of closest approach to the centre of the region of the proton-proton collision, referred to as *beamspot*, such that the contribution from pileup is reduced. The spatial resolution of the PV reaches values as low as $20\text{ }\mu\text{m}$ for vertices reconstructed with at least 50 tracks. In case more than one PV is reconstructed per event, the one for which the quantity $\sum p_T^2$, where the sum runs over the tracks associated with the vertex, is selected.

While the presence of at least one reconstructed PV in the event is required in the two searches presented in the thesis, displaced vertices also play an important role, either to identify jets originating from the hadronisation of a b quark, or to select events possibly containing a long-lived heavy neutrino.

2.4.2 Jets

Quarks and gluons produced in proton-proton collisions cannot be resolved individually as they go through the hadronisation process, introduced in Section 1.3, creating a cascade of particles collimated along the direction of flight of the original parton. Therefore, the experimental signature of a parton is a so-called jet. The energy and direction of the jet reflect that of the parton originating the hadronic shower. Jets are at the core of the supersymmetry analysis presented in Chapter 4.

Jets primarily consist of charged and neutral hadrons. The hadrons are clustered together using the anti- k_T algorithm [40] with a distance parameter equal to 0.4. This algorithm is both infrared- and collinear-safe, and reconstructs jets with a cone-like structure. The jet $\vec{p}_T$ is defined as the vectorial sum of the $\vec{p}_T$ of the particles clustered into the jet. To prevent the inclusion of particles originating from pileup interactions, which would lead to an overestimation of the jet energy, charged hadrons whose track is associated with a pileup vertex are discarded. Moreover, jets identified as originating from detector noise are rejected.

The measurement of the jet energy may be affected by various effects such as (i) detector

noise, (ii) inefficiency in the HCAL cluster reconstruction, (iii) non-uniform detector response, and (iv) leftover pileup contributions. Therefore, a series of corrections, referred to as *jet energy corrections* [41], is applied such that the averaged energy of jets corresponds to that of the underlying partons. Moreover, additional corrections are applied on simulated events such that the associated energy is in agreement with that obtained in data.

Finally, dedicated algorithms are designed to identify jets originating from the hadronisation of b quarks. Such jets are referred to as *b-jets* in the following. In the context of this work, the b-tagging algorithm used to identify b-jets is the DeepCSV algorithm [42]. It is a deep neural network that exploits typical signatures of beauty hadron decays to discriminate between b-jets and light-flavour jets, i.e. jets originating from the hadronisation of gluons or of u, d, or s quarks. Such signatures are for instance the presence of a decay vertex displaced with respect to the beamspot, and of particles with large values of the impact parameter of their associated tracks. These signatures follow from the fact that B hadrons have a relatively long lifetime. Another typical signature of the decay of beauty hadrons is the presence of non-isolated charged leptons, as B hadrons have a larger branching fraction to leptonic and semileptonic decays than the other hadrons.

2.4.3 Missing transverse momentum

When protons collide, the total longitudinal momentum of the system cannot be retrieved, as the system is boosted along the beam direction. Instead, the total transverse momentum is constrained to be zero, as the collision does not take place in the $x - y$ plane. This feature is exploited to compute the missing transverse momentum, $\vec{p}_T^{\text{miss}}$, defined such that the vectorial sum of the $\vec{p}_T$ of all the reconstructed particles with $\vec{p}_T^{\text{miss}}$ is null. Its magnitude is labelled as p_T^{miss} . The good hermiticity of the CMS detector is a key feature for a performant reconstruction of this object. Moreover, we note that the jet energy corrections introduced above are accounted for to allow for a robust computation of $\vec{p}_T^{\text{miss}}$.

While imperfect instrumentation and reconstruction algorithms may bias $\vec{p}_T^{\text{miss}}$ to non-zero values, genuine $\vec{p}_T^{\text{miss}}$ can be produced in the event by undetectable particles, such as SM neutrinos, or hypothetical new physics particles. For instance, the lightest supersymmetric particle in the search described in Chapter 4 gives rise to sizable values of p_T^{miss} .

2.4.4 Muons

Muons are reconstructed from the hits in the tracking system and in the muon spectrometer [43]. The hits in the tracker are fitted to reconstruct tracks, while these in the muon spectrometer are fitted to reconstruct so-called *segments*.

Two families of reconstructed muons are of interest in this work: the *global* muons and the *tracker* muons. The former muons are reconstructed starting from segments in the muon spectrometer that are then matched to a track in the tracking system. Instead, the latter muons are reconstructed starting from tracks in the tracking system that are extrapolated to the muon spectrometer and that must be compatible with at least one DT or one CSC segment of a muon

station.

The associated reconstruction algorithms fit the trajectory with the requirement that the track in the silicon tracker is consistent with originating from the PV. This requirement makes the search for long-lived heavy neutrinos challenging, as the muons in this search are assumed to originate from a displaced vertex rather than the primary one. A dedicated signal-to-background discriminator will be designed to enhance the sensitivity of the search, as discussed in Section 6.4.3, taking as input the impact parameters of the muon as well as the number of layers crossed by the particle in the pixel and strip detectors.

A useful quantity that is exploited in the two searches presented in the thesis is the PF isolation, I^{PF} , of the muon. It is defined as the scalar sum of the transverse momenta of the charged and neutral hadrons comprised within a cone of size ΔR around the muon, from which the contribution from pileup is subtracted. The quantity ΔR between two objects is defined as

$$\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}, \quad (2.2)$$

where $\Delta\eta$ and $\Delta\phi$ are the difference in pseudorapidity and in azimuthal angle of the two objects, respectively.

2.4.5 Electrons and photons

Electrons [44] and photons [45] are reconstructed from energy deposits in the ECAL crystals. Typically, the electromagnetic shower induced by an incoming particle spreads over several neighbouring crystals. The signals from these crystals are gathered into energy clusters, using the *clustering* algorithm. Additionally, as electrons and photons interact with the tracker material, they may convert, such that not a single particle, but a collection of particles, spread along the ϕ angle, reach the ECAL. The energy of the incoming particle is recovered by aggregating the different clusters into a single supercluster, thanks to the *superclustering* algorithm. An illustration of a supercluster is shown in Fig. 2.5.




Figure 2.5: Sketch of a supercluster.

While these algorithms are common to the reconstruction of both the electrons and photons, the electrons must additionally be reconstructed from the combination of a supercluster with a track in the silicon tracker. Moreover, the PF isolation, I^{PF} , defined above for the muons is also used to characterise the electrons and photons.

The clustering algorithm is therefore at the core of the electron and photon reconstruction. A part of the thesis was dedicated to the optimisation of its performance in view of the Run 3 of data taking, as described in the next section.




Figure 2.6: Average noise level, σ_{noise} , predicted in the ECAL barrel (left) and endcaps (right) for the Run 3 of data taking as a function of $|\eta|$. The different coloured lines correspond to different total accumulated luminosities. The figures are taken from Ref. [46].

2.5 Optimisation of the ECAL clustering algorithm

As the electromagnetic calorimeter is more and more exposed to radiation, the noise level inside the detector increases. As a consequence, the particle reconstruction algorithms must be optimised such that they achieve stable performance under the evolving conditions. An important project carried out during the PhD consisted in optimising the ECAL clustering algorithm, introduced in Section 2.4.5, in view of the Run 3 of data taking. The entire project, from the design of the optimisation strategy to the extraction of the final results, was part of the thesis work. It was conducted during the long shutdown that followed the end of the Run 2 of data taking and was performed in the context of the service task owed to the collaboration.

The predicted noise level, σ_{noise} , for different total accumulated luminosities is shown in Fig. 2.6 as a function of the absolute value of the pseudorapidity, $|\eta|$, in the ECAL barrel and endcaps. The noise level increases with the total accumulated luminosity as a consequence of the radiation (i) damaging the photosensors and (ii) reducing the transparency of the PbWO_4 scintillating crystals. Moreover, we verify that the noise level is higher in the endcaps than in the barrel, as the former part of the detector is more exposed to radiation.

As anticipated in Section 2.4.5, the clustering algorithm is an important step in the measurement of the energy of the electrons and photons. A complete description of the algorithm is presented in Ref. [25]. It consists in aggregating the energy deposited by the particle in neighbouring crystals into clusters. The algorithm is designed such that (i) the contribution from noise and pileup is reduced and that (ii) the case of overlapping particles is mitigated. The algorithm can be summarised in three main steps, illustrated in Fig. 2.7:

1. The algorithm takes as input the crystal energy deposits that are greater than the *crystal energy threshold*.
2. The seed of the cluster is identified as the crystal having the largest energy with respect to its neighbours, provided that this energy is greater than the *seeding energy threshold*.




Figure 2.7: Illustration of the main steps of the ECAL clustering algorithm.

3. The clusters are formed by gathering the neighbouring crystals around the seed and computing the energy fraction of a crystal to a given cluster to account for potential overlap between incoming particles. The energy of the cluster is then obtained as the weighted sum of the energy of the gathered crystals.

The application of the crystal and seeding energy thresholds aims to exclude the contribution from noise and pileup when forming the clusters. These thresholds are expressed as multipliers of σ_{noise} , such that their magnitude increases with larger values of $|\eta|$. For crystals in the barrel, the value of σ_{noise} is averaged over a ring of constant $|\eta|$, while for crystals in the endcaps, the value of σ_{noise} is taken per crystal. Indeed, as the latter region is more exposed to radiation, taking into account the state of the individual crystals was found to improve the overall performance of the algorithm.

The optimisation of the clustering algorithm consists in tuning the value of the crystal and seeding energy thresholds, by finding the optimal multipliers such that the noise is maximally rejected and the signal efficiency is preserved as much as possible. The tuning is performed using simulated events in which photons interact with the ECAL and are reconstructed through the clustering algorithm with modified parameters. This study is performed using σ_{noise} corresponding to a total integrated luminosity of 450 fb^{-1} , which is the luminosity predicted to be accumulated from the start of Run 1 to the end of Run 3. A total of 10 pairs of multipliers for the crystal and seeding thresholds are tested. The performance for each pair is assessed in different bins of energy and $|\eta|$ based on:

1. The energy resolution, defined as the relative width of the ratio of reconstructed to simulated energy, extracted from a maximum likelihood fit to the distribution of the ratio modelled with a double-sided Crystal Ball function [47, 48].
2. The signal efficiency, defined as the ratio of the number of reconstructed clusters matched to a simulated photon to the total number of simulated photons.
3. The noise rate, defined as the ratio of the number of clusters that are not matched to a simulated photon to the total number of clusters.




Figure 2.8: Energy resolution (*left*), signal efficiency (*middle*), and noise rate (*right*) as functions of the multipliers of σ_{noise} for the crystal and seeding thresholds in a representative energy and $|\eta|$ bin.

These quantities are shown in Fig. 2.8 as functions of the thresholds multipliers in a representative bin of energy and $|\eta|$. As expected, the lower is the seeding energy threshold, the larger is the efficiency and the lower is the noise rejection. On the other hand, the value of the crystal energy threshold mainly influences the resolution. The optimal value of the thresholds pair is the one that provides a trade-off between the antagonist trends.

The optimal crystal and seeding energy thresholds were found to be both equal to $3\sigma_{\text{noise}}$ for $|\eta| \leq 2.5$ and $4\sigma_{\text{noise}}$ for $|\eta| > 2.5$. Moreover, a dedicated study showed that the thresholds should be updated once per year, with σ_{noise} taken as the noise level predicted at mid-year.

Figure 2.9 shows as a function of $|\eta|$ the performance in the resolution, efficiency, and noise rate obtained with the tuned and reference thresholds for a detector with conditions corresponding to an integrated luminosity of 450 fb^{-1} . The reference thresholds correspond to these employed at the end of the Run 2 of data taking, which would have been used without the optimisation study. The resolution obtained with the tuned clustering algorithm is improved by up to 40% for $|\eta| > 1.85$. Because the tuned thresholds are tighter than the reference ones, a loss in efficiency of up to 45% is obtained for $|\eta| > 2.4$. However, this loss is limited compared to the reduction of the noise, which is almost entirely rejected across the entire $|\eta|$ range. Further figures showing the performance of the clustering tuning are presented in Appendix A.

This work was presented at the TIPP conference in 2021 [49]. An internal CMS note as well as all the computing framework were put at the disposal of the collaboration, such that the optimisation can be easily repeated in the future. All the CMS analyses making use of electrons, photons, jets, or missing transverse momentum with Run 3 data benefit from this work.




Figure 2.9: Energy resolution (*upper left*), signal efficiency (*upper right*), and noise rate (*lower*) obtained with the tuned and reference thresholds as a function of $|\eta|$. The performance is shown for simulated photons with transverse energy between 5 and 10 GeV that interact with the ECAL with conditions predicted for a total integrated luminosity of 450 fb^{-1} .

Part II

Search for new physics in the supersymmetric sector

Chapter 3

Main concepts of supersymmetry

Supersymmetry (SUSY) is a theoretical framework that extends the standard model with the postulate of an additional fundamental symmetry between fermionic and bosonic states. For each SM particle, the theory predicts the existence of a supersymmetric partner, known as a *superpartner*, whose spin differs by 1/2 unit. The supersymmetric operator, $\mathcal{Q}$, therefore transforms a bosonic state, $|b\rangle$, into a fermionic one, $|f\rangle$, and vice versa:

$$\mathcal{Q}|b\rangle = |f\rangle \quad \text{and} \quad \mathcal{Q}|f\rangle = |b\rangle. \quad (3.1)$$

The consequences of the introduction of this additional symmetry are discussed in Section 3.1 in the context of the minimal supersymmetric standard model (MSSM). In particular, we will motivate SUSY by that it provides, within a robust framework, answers to several of the limitations of the SM, discussed in Section 1.5. A step closer to the experimental work performed in the thesis is taken in Section 3.2, where the expected signatures from SUSY processes at hadron colliders are described. Finally, an overview of the signal models constrained in this work is given in Section 3.3.

The aim of this chapter is to provide the key concepts of supersymmetry, in view of giving context to the experimental search described in Chapter 4. A detailed overview of the theoretical building blocks and predictions of SUSY is for instance presented in Refs. [50, 51].

3.1 The minimal supersymmetric standard model

The minimal supersymmetric standard model is a supersymmetric extension of the SM through the introduction of a single SUSY operator. The states of a given SM particle and of its associated superpartner belong to the same irreducible representation of the SUSY algebra, called a *supermultiplet*. Because the fermionic and bosonic members of the supermultiplet are in the same representation of the gauge group, they have identical charges. As a consequence, the SM particle and its superpartner should be identical particles, except for their spin.

However, it is established that the particle and its superpartner do not have the same mass. Indeed, as the SM and SUSY particles have the same gauge interactions, with identical coupling strengths, the observation of SUSY particles would be experimentally straightforward if they

Table 3.1: Content of sparticles in the MSSM.

Type	Name	Spin	Gauge eigenstates
Sfermions	Squarks	0	$\tilde{u}_L, \tilde{u}_R, \tilde{d}_L, \tilde{d}_R$ $\tilde{s}_L, \tilde{s}_R, \tilde{c}_L, \tilde{c}_R$ $\tilde{t}_L, \tilde{t}_R, \tilde{b}_L, \tilde{b}_R$
	Sleptons	0	$\tilde{e}_L, \tilde{e}_R, \tilde{\nu}_e$ $\tilde{\mu}_L, \tilde{\mu}_R, \tilde{\nu}_\mu$ $\tilde{\tau}_L, \tilde{\tau}_R, \tilde{\nu}_\tau$
Gauginos	Winos	1/2	$\tilde{W}^\pm, \tilde{W}^0$
	Bino	1/2	$\tilde{B}^0$
	Gluinos	1/2	$\tilde{g}$
Higgsinos	Higgsinos	1/2	$\tilde{H}_u^+, \tilde{H}_u^0$ $\tilde{H}_d^-, \tilde{H}_d^0$

would have the same mass as their associated SM partner. Yet, no supersymmetric particles have been discovered so far. Therefore, SUSY cannot be an exact symmetry, and a mechanism for symmetry breaking must be introduced [51]. As the different breaking mechanisms that have been established up to now do not predict the value of the masses of the superpartners, these must be measured experimentally.

The introduction of supersymmetry to the model at least doubles the particle content of the universe, and adds more than a hundred new parameters to the theory. The supersymmetric particles, referred to as *sparticles*, can be classified into three categories: (i) the sfermions, which are the bosonic partners of the SM fermions; (ii) the gauginos, which are the fermionic partners of the SM vector bosons; (iii) and the higgsinos, which are the fermionic partners of the SM Higgs scalar bosons. After the spontaneous breaking of the electroweak symmetry, three neutral and two charged Higgs bosons are predicted in the MSSM.

The minimal content of sparticles in the MSSM is given in Table 3.1, together with the associated spin. For instance, the superpartners of the left-handed and right-handed top and bottom quarks are the left-handed and right-handed stop ($\tilde{t}_L$ and $\tilde{t}_R$) and sbottom ($\tilde{b}_L$ and $\tilde{b}_R$) particles, respectively; and the superpartner of the gluons are the gluinos ($\tilde{g}$). After the electroweak symmetry breaking, the neutral gauginos mix to give the zino ($\tilde{Z}$) and the photino ($\tilde{\gamma}$). Moreover, four mass eigenstates, referred to as the neutralinos ($\tilde{\chi}_{1-4}^0$), arise from the mixing of the neutral gauginos and the neutral Higgsinos. On the other hand, the mixing of their charged counterparts gives rise to another two mass eigenstates, referred to as the charginos ($\tilde{\chi}_{1-2}^\pm$). The list of neutralinos and charginos is given in Table 3.2. The mixing also takes place in the sfermion sector, but is only significant for the third generation of squarks and sleptons.

***R*-parity**

The *R*-parity is a multiplicative quantum number, defined for a given particle or sparticle as

$$P_R = (-1)^{3B+L+2S}, \quad (3.2)$$

Table 3.2: Neutralinos and charginos in the MSSM.

Name	Gauge eigenstates	Mass eigenstates
Neutralinos	$\tilde{B}^0, \tilde{W}^0, \tilde{H}_u^0, \tilde{H}_d^0$	$\tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0$
Charginos	$\tilde{W}^\pm, \tilde{H}_u^\pm, \tilde{H}_d^\pm$	$\tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm$

where B is the baryon number, L is the lepton number, and S is the spin of the particle or sparticle. The SM particles have even R -parity ($P_R = +1$), while the sparticles have odd R -parity ($P_R = -1$). This quantum number is predicted to be conserved in the MSSM, although it is not a strict requirement [52]. The conservation of the R -parity has the following important consequences:

1. Supersymmetric particles can only be produced in pairs in a process.
2. Supersymmetric particles can only decay into an odd number of sparticles.
3. The lightest supersymmetric particle (LSP) is stable.

As will be discussed in Section 4.1.1, these features will have important implications for the design of the strategy for the experimental search conducted in the thesis work.

Cosmological observations suggest that it is unlikely that the LSP, if stable, is electrically charged [53]. Within the MSSM, the LSP is therefore predicted to be a neutral particle, which only interacts via the weak and gravitational interactions. Possible candidates for the LSP are (i) the lightest neutralino, $\tilde{\chi}_1^0$, (ii) the lightest superpartner of neutrinos, and (iii) the superpartner of the spin-2 graviton, named the gravitino. The first candidate is the one considered in this work.

Solutions to open questions

The introduction of supersymmetry as an extension of the SM is particularly well motivated, as it provides solutions to some of the questions left open in the SM, discussed in Section 1.5.

First of all, a strong motivation for SUSY is that the hierarchy problem can be solved within this framework [54]. Indeed, in an equivalent manner as the corrections from fermionic loops, corrections to the Higgs boson mass from sfermionic loops must also be taken into account. These additional corrections cancel the overall quadratic divergencies, as (i) the fermions and sfermions couple to the Higgs boson with the same strength, and as (ii) the loop corrections from sfermions have the opposite sign with respect to that from fermions, because of the scalar nature of these sparticles. As a consequence, the level of fine-tuning is drastically reduced, and the naturalness of the Higgs boson mass is preserved.

Second of all, the MSSM may provide a candidate for dark matter [55]. As discussed above, a consequence of the R -parity conservation is the stability of the LSP. In case this particle is indeed electrically neutral and is stable for a longer time than the age of the universe, it may be dark matter. The most promising candidate is the lightest neutralino, $\tilde{\chi}_1^0$ [50].




Figure 3.1: Evolution of the inverse gauge couplings associated with the U(1), SU(2), and SU(3) gauge symmetries, for the case in which the sparticles have masses in the TeV range. The figure is taken from Ref. [1].

Another motivation for the existence of supersymmetry is the unification of the coupling constants [56]. Indeed, these parameters vary as a function of the energy scale. As opposed to the SM, the MSSM predicts that the strong, electromagnetic, and weak gauge couplings unify at an energy scale of about 10^{16} GeV. This behaviour is illustrated in Fig. 3.1. Although this feature does not strictly speaking solve an issue with the SM, its elegance provides yet another motivation for the existence of supersymmetry.

In the same spirit, the prediction of the mass of the Higgs boson within the MSSM increases the motivation for SUSY searches. Indeed, within this model, the mass of the scalar Higgs boson is predicted to be bounded from above by about 130 GeV, when including the radiative corrections from top quark and squark loops [57]. The discovery in 2012 of a boson with a mass of about 125 GeV compatible with the Higgs properties therefore contributed to motivating the large effort deployed to find SUSY at the LHC.

3.2 Probing supersymmetry at hadron colliders

The strong predictions from supersymmetry, together with the simplicity of the principle on which it is constructed, have led it to become one of the most promising theories of new physics. A huge experimental effort was thus gathered in the past decades to discover it. Early searches were performed at experiments placed at (i) LEP, which is an electron-positron collider that reached a maximal centre-of-mass energy $\sqrt{s}_{\text{max}} = 209$ GeV; at (ii) HERA, which is an electron-proton collider with $\sqrt{s}_{\text{max}} = 318$ GeV; and at (iii) the Tevatron, which is a proton-antiproton collider with $\sqrt{s}_{\text{max}} = 1.96$ TeV. No sign for SUSY was found in any of the data analysed. A detailed review of the corresponding results is presented in Ref. [58]. With the start of the operation of the LHC in 2009, prospects to discover SUSY were optimistic, as the higher centre-of-mass energy of the collider allows for probing sparticles in higher mass regimes. Yet, no significant evidence for supersymmetry has been found so far.




Figure 3.2: Feynman diagrams for gluino (upper) and squark (lower) production via gluon-gluon fusion (left) and quark-quark annihilation (right).

Probing SUSY at hadron colliders, in particular at the LHC, is as promising as it is challenging. On the one hand, the delivered energy allows for searching for sparticles at high masses. Moreover, the hadron collisions make the search for strong SUSY processes possible, while also giving the chance to probe rarer SUSY processes. On the other hand, events at hadron colliders tend to have complex signatures, as (i) final states include fragments of the colliding hadrons that do not participate in the hard reaction, as (ii) sparticles cascade decays may be deeper, and as (iii) the SM backgrounds are copiously produced, and must be handled such as not to mask the signatures of sparticles decays in the data.

In the following, we discuss the main sparticle production and decay signatures, as well as the main backgrounds expected at a hadron collider. A particular focus is given on processes with strong SUSY production and R -parity conservation, as they are the class of processes on which the direct search presented in Chapter 4 is based.

The largest SUSY cross sections at the LHC come from strong processes, such as

$$gg \rightarrow \tilde{g}\tilde{g}, \quad gg \rightarrow \tilde{q}_i\tilde{q}_j^*, \quad gq \rightarrow \tilde{g}\tilde{q}_i, \quad q\bar{q} \rightarrow \tilde{g}\tilde{g}, \quad q\bar{q} \rightarrow \tilde{q}_i\tilde{q}_j^*, \quad \text{and} \quad qq \rightarrow \tilde{q}_i\tilde{q}_j. \quad (3.3)$$

Examples of Feynman diagrams are shown in Fig. 3.2. The cross sections for pair-production of gluinos, light-flavour squarks, and heavy-flavour squarks are shown in Fig. 3.3 as a function of the produced sparticle mass, M . These cross sections are obtained at the NNLO_{Approx}+NNLL accuracy for proton-proton collisions at $\sqrt{s} = 13$ TeV [59]. At a given mass M , the cross sections for electroweak processes are up to two orders of magnitude smaller than those for strong processes [60]. This makes a search in the strong sector one of the most promising channels to discover SUSY.

Signatures of a strong SUSY process with R -parity conservation at a hadron collider are

1. High energetic jets, as a consequence of the hadronisation of the quarks produced in the decay of the sparticles.
2. A sizable missing transverse momentum, p_T^{miss} , defined in Section 2.4, following from the presence of the LSPs, which are undetectable. As the sparticles are pair-produced, there must be an even number of LSPs in the final state.




Figure 3.3: Cross section as a function of the sparticle mass, M , for gluino, light-flavour squark, and heavy-flavour squark pair-production in proton-proton collisions at $\sqrt{s} = 13$ TeV, computed at NNLO_{Approx}+NNLL accuracy. The width of the curves represents the associated theoretical uncertainties. The light-flavour squarks, $\tilde{u}$, $\tilde{d}$, $\tilde{s}$, and $\tilde{c}$, are assumed to be mass degenerate and the associated cross section is the sum between the left- and right-handed chiral states.

When decaying, sparticles produce cascade reactions inside the detector. Typically, final states originating from hadron collisions consist of N_ℓ leptons + N_j jets + p_T^{miss} , where either the lepton multiplicity, N_ℓ , or jet multiplicity, N_j , might be null.

Several SM processes yield similar final states as the new physics processes. In particular, an important contribution to the background at a hadron collider arises from QCD processes. Strategies must therefore be put in place to reduce this overwhelming source of background and be able to disentangle the signatures of SUSY processes in the data. For instance, imposing a minimal threshold on p_T^{miss} allows for a significant reduction of this background component. More sophisticated strategies employ kinematic variables designed to offer a good discrimination between new physics and background events. The M_{T2} variable, on which the search performed during the thesis is based, is one of such examples.

Other than QCD processes, significant contributions to the background come from Z+jets and W+jets processes, in which the bosons decay into neutrinos, and $t\bar{t}$ +jets processes. As the latter two processes can typically give rise to leptons in the final state, an efficient way of reducing these background sources consists in imposing $N_\ell = 0$, and thus in performing an all-hadronic search. This is the strategy employed for the SUSY search described in the next chapter.

3.3 Simplified supersymmetric models

The MSSM has over one hundred new parameters compared to the SM, making the interpretation of the data complex. Simplified SUSY models [61, 62] were derived to facilitate the interpretation




Figure 3.4: Feynman diagrams representing the gluino-mediated processes constrained in this work. The gluino decays into a light flavour (*left*), beauty (*middle*), and top (*right*) quarks are shown.




Figure 3.5: Feynman diagrams representing the squark-mediated processes constrained in this work. The direct production of light flavour (*left*), beauty (*middle*), and top (*right*) squark pairs is shown.

of the results. These models constrain the parameter space of the MSSM by imposing the following conditions:

1. Each model represents a specific SUSY process, i.e. the pair-production of one type of particles that decay with a specific decay mode.
2. The particles that are not involved in the process are assumed to be heavy enough such that they decouple from the new physics process and hence do not influence it.
3. The branching fraction of the decay mode of the particles under scrutiny is assumed to be 100%.

These conditions reduce the number of free parameters to two, namely the mass of the sparticle and the mass of the LSP. Although the simplified SUSY models do not reflect the wide range of variations of the MSSM, there are useful for benchmarking purposes. Indeed, they can be used to assess the sensitivity of a search, and hence compare the sensitivity reach of different analyses. Moreover, they would allow for a first characterisation of the new physics signal, in case of a discovery.

In the context of this work, six different simplified models of R -parity conserving SUSY involving coloured particles are constrained by the data. They belong to two categories: the gluino-mediated decays, shown in Fig. 3.4, and the squark-mediated decays, shown in Fig. 3.5. For each category, three decay modes involving either light-flavour quark(s), beauty quark(s), or top quark(s) are considered. The quarks are produced together with the LSP, taken to be

Table 3.3: List of the three gluino-mediated and three squark-mediated SUSY models considered in this work. For each signal model, the sparticle production and decay modes are reported.

	Sparticle production	Sparticle decay
Gluino-mediated decays	$pp \rightarrow \tilde{g}\tilde{g}$	$\tilde{g} \rightarrow q\bar{q}\tilde{\chi}_1^0$
	$pp \rightarrow \tilde{g}\tilde{g}$	$\tilde{g} \rightarrow b\bar{b}\tilde{\chi}_1^0$
	$pp \rightarrow \tilde{g}\tilde{g}$	$\tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$
Squark-mediated decays	$pp \rightarrow \tilde{q}\tilde{q}$	$\tilde{q} \rightarrow q\tilde{\chi}_1^0$
	$pp \rightarrow \tilde{b}\tilde{b}$	$\tilde{b} \rightarrow b\tilde{\chi}_1^0$
	$pp \rightarrow \tilde{t}\tilde{t}$	$\tilde{t} \rightarrow t\tilde{\chi}_1^0$

the lightest neutralino, $\tilde{\chi}_1^0$. The light-flavour quarks correspond to the quarks of the first two generations, i.e. the up, down, strange, and charm quarks. For the squark-mediated signal to light-flavour quarks, the squarks are considered to be mass degenerate. The full list of probed signal models is given in Table 3.3.

Chapter 4

Search for supersymmetry using the M_{T2} variable

A search for supersymmetry is performed in all-hadronic final states with large missing transverse momentum using proton-proton collision data collected at a centre-of-mass energy $\sqrt{s} = 13$ TeV. The analysis exploits the Run 2 data sets, collected in 2016, 2017, and 2018, corresponding to a total integrated luminosity of 137 fb^{-1} .

The search probes supersymmetric processes produced via the strong interaction and conserving R -parity. As discussed in Section 3.1, such processes involve the pair-production of sparticles, either squarks or gluinos, that decay promptly into quarks and the LSP, predicted to be a neutral and stable particle. The hadronisation of the quarks or gluons gives rise to a large hadronic activity in the final state, while the presence of the LSP, which remains undetected, can manifest itself through a sizable transverse momentum imbalance in the event.

Performing a search in the all-hadronic channel has the advantage to reduce the electroweak SM backgrounds associated with the presence of leptons in the final state, as discussed in Section 3.2. On the other hand, the M_{T2} variable, defined in Section 4.1.1, is used to reduce the large contribution from QCD SM backgrounds. Furthermore, because this quantity is a powerful discriminator between signal and SM background processes, it is used as the discovery variable in the search.

After applying the event reconstruction and selection criteria, described in Sections 4.3 and 4.4, respectively, the number of observed data events is compared to the number of predicted events from SM background processes. Such a comparison is performed in multiple event categories, defined in Section 4.5. A significant excess of data events over the SM prediction, primarily expected for large M_{T2} values, as discussed in Section 4.6, would be a sign for the presence of new physics in the data. As detailed in Section 4.7, background events are expected to arise from three main classes of processes:

1. $Z \rightarrow \nu\bar{\nu}$ background: Z +jets events where the Z boson decays invisibly into neutrinos. This background source is irreducible and is the dominant one in most of the event categories.
2. Lost-lepton background: events involving a leptonic W boson decay, for which the charged

Table 4.1: List of previous iterations of the search together with the reference to their respective publication. The searches were performed using proton-proton collision data collected at $\sqrt{s} = 7, 8,$ or 13 TeV and corresponding to different integrated luminosities.

$\sqrt{s}$ (TeV)	Integrated luminosity (fb^{-1})	Year of publication	Reference
7	5	2012	[63]
8	20	2015	[64]
13	2	2016	[65]
13	36	2017	[66]

lepton is either out of acceptance, not reconstructed, not identified, or not isolated. This background may arise from W +jets and $t\bar{t}$ +jets processes. The genuine missing transverse momentum from the neutrino makes of this background the second largest contribution.

3. QCD multijet background: events with significant missing transverse momentum arising from jet mismeasurement or detector noise. The contribution of this background is considerably reduced by the application of a minimum threshold on the value of the M_{T2} variable, as described in Section 4.1.1.

Simultaneous maximum likelihood fits to the data in all the event categories are finally used to extract constraints on the parameter space of the simplified supersymmetric models, defined in Section 3.3.

The present search was done in collaboration with members of U.S. institutes at UCSB and UCSD. It builds on the work established in previous iterations of the analysis [63–66], listed in Table 4.1, which were performed at different centre-of-mass energies and exploited data sets corresponding to different integrated luminosities.

The aim of the thesis work was to repeat the analysis, from the sample production to the extraction of the results, using the complete data sample collected during the Run 2 of the LHC, which corresponds to an increase in the luminosity of about a factor 4 with respect to the previous iteration of the analysis [66]. In particular, a great effort was invested in the monitoring of the quality of the data and of the simulated events. Moreover, a significant part of the work was invested in the study and validation of the estimation methods of the $Z \rightarrow \nu\bar{\nu}$ and lost-lepton backgrounds. Finally, a key aspect of the work was the combination of the three years of data taking to extract the final results.

The outcome of the search was presented at the Moriond/EW 2019 conference [67] and was published in 2020 in the *European Physical Journal C* [68].

This chapter is organised as follows: the definition and properties of the M_{T2} variable and other useful kinematic variables are presented in Section 4.1. A description of the data samples is then given in Section 4.2. Sections 4.3, 4.4, and 4.5 describe the event reconstruction, selection, and categories, respectively. The expected signal yields for the probed signal models are obtained following the description given in Section 4.6, and the three background sources are discussed in Section 4.7. Finally, the results and their interpretation are presented in Section 4.8, and the main results are summarised in Section 4.9.

4.1 Search variables

4.1.1 The M_{T2} variable

When searching for a new particle, the reconstruction of its mass often plays a key role and may, in some instances, be the discovery variable that is used to detect its presence. In the case in which the sought-after particle, produced at a hadron collider, decays semi-invisibly into a set of detectable particles plus an undetectable particle, the reconstruction of the total invariant mass is not possible, as the longitudinal component of the momenta cannot be measured. An example of such a semi-invisible decay is illustrated in Fig. 4.1 (left) for the process $W^+ \rightarrow e^+ \nu_e$, where the positron, e^+ , is visible and the neutrino, ν_e , is invisible. For such decays, one can nonetheless reconstruct the transverse mass, m_T , which is computed using only the projections of the final-state momenta onto the plane transverse to the beam direction:

$$m_T^2 \equiv m_{\text{vis}}^2 + m_\chi^2 + 2(E_T^{\text{vis}} E_T^\chi - \vec{p}_T^{\text{vis}} \cdot \vec{p}_T^\chi), \quad (4.1)$$

where m_{vis} and m_χ are the masses of the visible and invisible systems, respectively; $\vec{p}_T^{\text{vis}}$ and $\vec{p}_T^\chi$ are the respective transverse momenta; and E_T^{vis} and E_T^χ are the respective transverse energies, defined, for a particle p , as

$$E_T^p \equiv \sqrt{m_p^2 + |\vec{p}_T^p|^2}. \quad (4.2)$$

By construction, the upper bound on the transverse mass is given by the mass of the semi-invisibly decaying mother particle. The transverse mass played a key role in the discovery of the W boson [69, 70].

The M_{T2} variable [71] is the generalisation of the transverse mass to processes in which a hadronic collision produces not a single, but a *pair* of massive particles that decay predominantly into one visible and one invisible particle. An example of such a decay is shown in Fig. 4.1 (right) for the R -parity conserving process with pair-produced squarks, $\tilde{q}$, that decay into an SM quark, q , and a stable neutralino, $\tilde{\chi}_1^0$. A further complication associated with processes involving pair-produced particles is that only the total missing transverse momentum,

$$\vec{p}_T^{\text{miss}} = \vec{p}_T^{\chi(1)} + \vec{p}_T^{\chi(2)}, \quad (4.3)$$

can be measured experimentally, as the individual transverse momenta of the two invisible systems, $\vec{p}_T^{\chi(1)}$ and $\vec{p}_T^{\chi(2)}$, cannot be resolved. This fact motivates the introduction of the M_{T2} variable, defined as

$$M_{T2}^2(m_\chi) \equiv \min_{\vec{p}_1 + \vec{p}_2 = \vec{p}_T^{\text{miss}}} \left\{ \max \left[m_T^{(1)}, m_T^{(2)} \right] \right\}, \quad (4.4)$$

where

- It is assumed that the invisible particles in the two decay chains, indexed with $i = 1, 2$, have the same mass, m_χ , i.e. $m_T^{(i)} = m_T^{(i)}(m_{\text{vis}(i)}, m_\chi, \vec{p}_T^{\text{vis}(i)}, \vec{p}_i)$, where $m_T^{(i)}$ is defined in Eq. (4.1).
- The maximal value between $m_T^{(1)}$ and $m_T^{(2)}$ is chosen, as these transverse masses would be bounded from above by the mass of the mother particle in case the mass, m_χ , of the invisible particles and their respective transverse momenta were known.




Figure 4.1: Schematic representation of events showing the SM decay $pp \rightarrow W^+$; $W^+ \rightarrow e^+ \nu_e$ (left), and the R -parity conserving SUSY decay $pp \rightarrow \tilde{q}\tilde{q}$; $\tilde{q} \rightarrow q\tilde{\chi}_1^0$ (right). The red and green colours highlight the visible and invisible final states, respectively.

- The minimisation over all possible combinations of momenta $\vec{p}_1$ and $\vec{p}_2$ whose sum gives the observed $\vec{p}_T^{\text{miss}}$, defined in Eq. (4.4), is performed to ensure that the M_{T2} value does not exceed the mass of the mother particle.

The mass m_χ is unknown and therefore appears as a free parameter in Eq. (4.4). In this work, M_{T2} is computed for $m_\chi = 0$, and it is implied in the following that $M_{T2} \equiv M_{T2}(0)$. Moreover, the visible particles are from now on considered to be quarks, which hadronise into jets in the detector.

To compute the M_{T2} variable, two further questions must be addressed. The first question concerns the generalisation of the M_{T2} variable to scenarios in which there are more than one visible particle per decay chain. Such scenarios may arise in events with initial state radiation (ISR) or pileup jets, but may also follow from the prediction of new physics models. For instance, the gluino-mediated decays, introduced in Section 3.3, give rise to two visible quarks per sparticle decay chain. In the context of this work, for events with more than two jets in the final state, the transverse momenta, $\vec{p}_T^{\text{vis}(1)}$ and $\vec{p}_T^{\text{vis}(2)}$, of the visible systems are not taken from the individual jets but rather from two groups of jets, referred to as *pseudojets*. The two pseudojets are reconstructed with the hemisphere algorithm described in Section 13.4 of Ref. [72]. This algorithm exploits the fact that the pair-produced particles may give rise to visible objects falling in distinct regions of the detector, reconstructed as hemispheres. The two pseudojets used as the visible systems in Eq. (4.4) are defined as the sum of the four-momenta of all the visible objects assigned to one or the other hemisphere.

The second question concerns the definition of the visible masses, $m_{\text{vis}(1)}$ and $m_{\text{vis}(2)}$, in Eq. (4.4). While assigning the mass of the pseudojets would appear as a natural choice, it was studied in previous iterations of the analysis [63, 64] that this choice was suboptimal, as mismeasured multijet background events with large pseudojets masses may give rise to large values of M_{T2} . This situation is unwanted because the tail of the M_{T2} distribution is the region with the greatest sensitivity to new physics topologies, as discussed in the following. Instead, it was found that using massless pseudojets keeps the value of the M_{T2} small for these background




Figure 4.2: Distribution of the M_{T2} variable shown for the simulated SM backgrounds (stacked histograms) and a representative signal (coloured line). The signal represents the pair-production of gluinos of mass $M = 1500$ GeV that decay into two beauty quarks and the lightest neutralino of mass $m_{\tilde{\chi}_1^0} = 100$ GeV. The normalisation of the signal is increased by a factor 10 to improve the readability. The figure was provided by UCSB.

events. Therefore, the M_{T2} variable is computed in this work with $m_{\text{vis}(1)} = m_{\text{vis}(2)} = 0$.

While the M_{T2} variable was originally introduced to measure the mass of pair-produced particles decaying semi-invisibly, this variable was also found to serve as an efficient discriminant between new physics and SM background processes. To illustrate its good discrimination power, Fig. 4.2 shows the M_{T2} distribution for QCD and electroweak background processes as well as for a signal process involving the pair-production of gluinos decaying into beauty quarks and the lightest neutralino, $\tilde{\chi}_1^0$. We observe that the QCD production of multijet events is the largest contribution at small M_{T2} values, while the signal populates the tails of the distribution. To understand this behaviour, let us consider the analytical expression of the M_{T2} variable, for the case in which there are no ISR jets and where the visible and invisible systems are massless:

$$M_{T2}^2 = 2p_T^{\text{vis}(1)} p_T^{\text{vis}(2)} (1 + \cos \theta_{12}), \quad (4.5)$$

with θ_{12} the angle between the two pseudojets in the transverse plane. The value of M_{T2} is therefore maximal when the pseudojets are parallel and minimal when they are produced back-to-back. As new physics events contain genuine $\vec{p}_T^{\text{miss}}$, the final states tend to have acoplanar visible systems and the obtained values of M_{T2} are typically large. Instead, for QCD multijet background events, in which the $\vec{p}_T^{\text{miss}}$ originates from jet mismeasurement, it is likely that the pseudojets are reconstructed back-to-back, giving rise to small values of M_{T2} . Regarding electroweak background events, the obtained values of M_{T2} tend to be larger than for QCD multijet background events, as these background processes contain genuine $\vec{p}_T^{\text{miss}}$ from the neutrinos.

The good discrimination power of the M_{T2} variable between signal and SM background processes makes it a promising tool to detect the presence of supersymmetric particles in the

data. The present work thus uses M_{T2} as the discovery variable. In particular, it is designed as a search for an excess of observed events with respect to the SM prediction in the tails of the M_{T2} distribution.

4.1.2 Other kinematic variables

Some other kinematic variables are useful to characterise the main signatures of the probed final states, by quantifying the hadronic activity and the transverse momentum imbalance in the event. The variables relevant for this search are

- H_T , defined as the scalar p_T sum of all selected jets in the event, $H_T \equiv \sum_{i \in \{\text{selected jets}\}} p_T(i)$.
- $\vec{H}_T^{\text{miss}}$, defined as the negative of the vector p_T sum of all selected jets in the event, $\vec{H}_T^{\text{miss}} \equiv -\sum_{i \in \{\text{selected jets}\}} \vec{p}_T(i)$. Its magnitude is noted as H_T^{miss} .
- $\vec{p}_T^{\text{miss}}$, defined as the negative of the vector p_T sum of all reconstructed PF candidates in the event, $\vec{p}_T^{\text{miss}} \equiv -\sum_{i \in \{\text{PF candidates}\}} \vec{p}_T(i)$. Its magnitude is noted as p_T^{miss} .
- N_j , defined as the number of selected jets in the event.
- N_b , defined as the number of selected jets originating from the hadronisation of a b quark, hereafter referred to as b-jets, in the event.

4.2 Event samples

4.2.1 Data samples

Luminosity

The search is performed using proton-proton collision data, collected during 2016, 2017, and 2018 at a centre-of-mass energy of 13 TeV. The luminosity integrated over each year is listed in Table 4.2 and corresponds to a total of 137 fb^{-1} .

Table 4.2: List of integrated luminosity per year of data taking.

Year	Luminosity (fb^{-1})
2016	36
2017	41
2018	60
Total	137

Triggers for signal regions

The analysis probes final states with large hadronic activity and large transverse momentum imbalance. The data are selected using triggers with different thresholds on H_T , jet p_T , p_T^{miss} , and/or H_T^{miss} , listed in Table 4.3 for the different years of data taking. A logical OR between the

Table 4.3: List of hadronic, invisible, and mixed triggers used during the three years of data taking. For each trigger, the HLT thresholds on the value of H_T , p_T of any jet, p_T^{miss} , and H_T^{miss} are reported. The dash symbol “–” indicates that there is no threshold on the corresponding quantity.

Year	H_T (GeV)	Jet p_T (GeV)	p_T^{miss} (GeV)	H_T^{miss} (GeV)	Type
2016	> 900	–	–	–	hadronic
2016	–	> 450	–	–	hadronic
2016	–	–	> 120	> 120	invisible
2016	> 300	–	> 110	–	mixed
2017 & 2018	> 1050	–	–	–	hadronic
2017 & 2018	–	> 500	–	–	hadronic
2017 & 2018	–	–	> 120	> 120	invisible
2017 & 2018	> 800	–	> 75	> 75	mixed
2017 & 2018	> 500	–	> 100	> 100	mixed
2017 & 2018	> 60	–	> 120	> 120	mixed

different trigger lines is applied. The triggers select events based on requirements on (i) hadronic objects only (hadronic triggers), on (ii) invisible objects only (invisible triggers), or on (iii) both hadronic and invisible objects (mixed triggers). The hadronic triggers have high thresholds on H_T (at least > 900 GeV) to make the search in final states with large hadronic activity possible. Instead, invisible triggers are useful to extend the analysis to a kinematic region with lower values of H_T . Finally, the mixed triggers overlap with the hadronic and invisible triggers to cover for their potential inefficiency in the intermediate H_T region. The value of the thresholds used for the trigger is tightened at the level of the event selection, as discussed in Section 4.4, to ensure that the trigger efficiency is maximal ($> 99\%$).

Triggers for control regions

The different background sources are estimated using events in background-enriched control regions. Further details on the definition of the control regions and the corresponding background estimation methods are given in Section 4.7. The different control regions and their respective triggers may be summarised as follows

- Single-lepton (e/μ) control regions: events are selected using the same triggers as for the signal regions.
- Same-flavour dilepton ($ee/\mu\mu$) control regions: because the computation of H_T and p_T^{miss} is different in these control regions than in the signal regions, as discussed in Section 4.7, the same triggers as for the signal regions cannot be employed. Instead, events are selected using dielectron and dimuon triggers, with the requirement that the charged leptons are isolated, as well as single-muon and single-photon triggers, to recover potential inefficiencies of the dilepton triggers.
- Different-flavour dilepton ($e\mu$) control regions: for the same reasons as above, events are collected with different triggers as those for the signal regions. Events are selected using

Table 4.4: List of simulated background and signal samples. For each sample, the frameworks used for the event generation and the simulation of the detector response are reported.

Samples	Event generator	Event simulation
W+jets, Z+jets, $t\bar{t}$ +jets, multijet	MADGRAPH5 (LO)	GEANT4
$t\bar{t}$ +W, $t\bar{t}$ +Z	MADGRAPH5 (LO)	GEANT4
single top	MADGRAPH5 (NLO)	GEANT4
SUSY signal	MADGRAPH5 (LO)	GEANT4
2016		
2017 + 2018	MADGRAPH5 (LO)	CMS fast simulation

a combination of isolated and non-isolated $e\mu$ triggers, non-isolated single-muon triggers, and a single-photon trigger.

4.2.2 Simulated samples

Monte Carlo (MC) simulation of background and signal events is used for the design of the search, the estimation of the SM backgrounds, and the measurement of the efficiency of the new physics models.

Samples are simulated for the main background processes, W+jets, Z+jets, $t\bar{t}$ +jets, and multijet, for the secondary background processes, $t\bar{t}$ W and $t\bar{t}$ Z, and for each of the simplified SUSY models introduced in Section 3.3. For each signal model, samples are generated for gluino or squark masses within $300 < M < 3000$ GeV and LSP masses within $0 < m_{\tilde{\chi}_1^0} < 1500$ GeV, with a step size of 50 GeV.

The programs used for the event generation and the simulation of the detector response are listed in Table 4.4 for each simulated sample. The physics events are generated with either leading order (LO) or next-to-leading order (NLO) accuracy using the MADGRAPH5_aMC@NLO 2 generator [73], which is interfaced with PYTHIA 8.2 [74] for the fragmentation and parton showering. The detector response is then simulated using either the GEANT4 program [75] or the CMS fast simulation framework [76]. The simulated events are finally reconstructed using the same chain of algorithms as for collision data. The processing of the samples was performed by the responsible group at CMS and was not part of the thesis work.

A series of corrections is applied to the simulated samples as event weights, to correct the modelling of certain simulated quantities. These corrections are lepton scale factors, correcting for the lepton identification efficiencies; b-tagging scale factors, correcting for the b-jet tagging algorithm efficiency; pileup weights, correcting for the distribution of the number of vertices; and ISR weights, correcting for the distribution of ISR jets. The respective corrections were provided by the responsible groups at CMS, together with their uncertainties.

4.3 Event reconstruction

In the present search, final states with jets and p_T^{miss} are probed. An efficient and robust reconstruction and identification of these objects is therefore placed at the core of the analysis work. Moreover, charged leptons must also be reconstructed with two aims:

1. Veto the leptonic events.
2. Estimate the leptonic background sources. As discussed in Section 4.7, the estimation of these processes is based on dedicated single-lepton and dilepton control regions.

The reconstruction of jets, p_T^{miss} , and charged leptons was discussed in Section 2.4, where tracks and vertices were also defined. In the following, additional conditions on the objects kinematics, isolation, and identification specific to this analysis are presented.

Among the identification criteria applied to some objects, a condition on the displacement of the tracks may be applied, in the aim of selecting prompt processes, while rejecting contributions from pileup. The displacement of tracks is measured as the transverse and longitudinal impact parameters, $|d_{xy}|$ and $|d_z|$, computed with respect to the primary vertex (PV). Among the vertices reconstructed in the event, the PV is defined as the event vertex with the largest sum of square of the p_T of its tracks, provided that it is at most 25 cm away from the beamspot along the direction of the beam line.

Jets and p_T^{miss}

The reconstructed jets must satisfy $p_T > 30$ GeV and $|\eta| < 2.4$ as well as the identification criteria summarised in Table 4.5. Among the set of reconstructed jets, those that originate from the hadronisation of a b quark are identified using the DeepCSV algorithm [42], introduced in Section 2.4.2. The b-jets are selected with a cut on the output of this deep neural network such that the b-tagging efficiency is about 55–70% for jets with $20 < p_T < 400$ GeV, and that the misidentification rate is about 1–2% (10–15%) for light-flavour or gluon jets (c-jets). Moreover, energy corrections, introduced in the aforementioned section, are applied to all jets to account for the effect of non-uniform detector responses and for the contribution of pileup. The missing transverse momentum, p_T^{miss} , is reconstructed following the approach described in Chapter 2.4.3. Moreover, it is computed using the same corrections as for the jets.

Table 4.5: Jet identification criteria in monojet ($N_j = 1$) and multijet ($N_j > 1$) events for the different data sets. The dash symbol “–” indicates that no condition on the corresponding quantity is applied.

	$N_j = 1$	$N_j > 1$	
	2016-2018	2016	2017-2018
Charged hadron fraction	> 0.05	> 0	> 0
Charged electromagnetic fraction	< 0.99	< 0.99	–
Charged multiplicity	> 0	> 0	> 0
Neutral hadron fraction	< 0.80	< 0.99	< 0.90
Neutral electromagnetic fraction	< 0.70	< 0.99	< 0.90
Number of constituents	> 1	> 1	> 1

Leptons

Muon candidates are reconstructed either as *global* or *tracker* muons, defined in Section 2.4.4. They must satisfy $p_T > 10$ GeV, $|\eta| < 2.4$, $|d_z| < 0.5$ cm, and $|d_{xy}| < 0.2$ cm, and are required to be isolated. The PF isolation, I^{PF} , defined in Section 2.4.4, is computed for tracks within a cone of size $\Delta R = 0.2, 0.1$, or 0.05 for $p_T < 50$ GeV, $50 < p_T < 200$ GeV, or $p_T > 200$ GeV, respectively. Candidates with a ratio of the PF isolation to the reconstructed muon p_T , I^{PF}/p_T , greater than 0.2 are rejected.

Electrons are reconstructed with the algorithms described in Section 2.4.5, and must have $p_T > 10$ GeV and $|\eta| < 2.4$. Moreover, they must satisfy $I^{\text{PF}}/p_T > 0.1$. While the muons used to define either the veto or the background-enriched control regions are the same, electrons used to define the veto have looser identification criteria on the shower shape than those used to define the control regions.

Isolated tracks

For the sake of maximally rejecting leptonic events, a veto on charged PF candidates, identified as isolated tracks, is further applied. These tracks typically originate from charged leptons that do not pass all the identification criteria to be reconstructed as such, or from charged hadrons. Therefore, this additional veto improves the rejection of events with electrons, muons, or hadronically decaying tau leptons.

Different sets of requirements are applied depending on the type of the charged PF candidate. For the PF electrons and muons, tracks must satisfy $p_T > 5$ GeV, $|\eta| < 2.4$, $|d_z| < 0.1$ cm, and $|d_{xy}| < 0.2$ cm. They must also obey the isolation criterion $I^{\text{track}}/p_T^{\text{track}} < 0.2$, where I^{track} is defined as the sum of the p_T of all charged PF candidates satisfying $|d_z| < 0.1$ cm, within a cone of size $\Delta R = 0.3$, and where p_T^{track} is the transverse momentum of the isolated track. For the charged PF hadrons, the tracks must have $p_T > 10$ GeV, $|\eta| < 2.4$, $|d_z| < 0.1$ cm, $|d_{xy}| < 0.2$ cm, and a track isolation $I^{\text{track}}/p_T^{\text{track}} < 0.1$.

4.4 Event selection

In addition to the trigger and object selections, discussed in Sections 4.2.1 and 4.3, respectively, further requirements are applied on relevant kinematic variables to preferentially select signal-like events for the search, while suppressing SM background. These requirements are

- $N_j \geq 1$ and $H_T > 250$ GeV: these cuts select events with a sizable hadronic activity.
- $p_T^{\text{miss}} > 250$ GeV for $H_T < 1200$ GeV, otherwise $p_T^{\text{miss}} > 30$ GeV: these requirements are applied to make sure that the trigger efficiency is at the plateau. At large H_T , a minimal cut on p_T^{miss} is applied for a sound computation of the M_{T2} variable.
- For events with at least two jets, $M_{T2} > 200$ GeV for $H_T < 1500$ GeV, otherwise $M_{T2} > 400$ GeV: as discussed in Section 4.1.1, imposing a minimum threshold on the value of

Table 4.6: List of final state objects and corresponding search regions that trigger the additional HCAL veto. Events in the 2018 data sets impacted by the two HCAL modules failure are rejected if they contain at least one of these objects in the corresponding search region and with $-4.7 < \eta < -1.4$ and $-1.6 < \phi < -0.8$.

Object	Search region
Jet with $p_T > 20$ GeV	$N_j = 1$ or $H_T < 575$ GeV
Jet with $p_T > 30$ GeV	$N_j \geq 2$ and $H_T > 575$ GeV
Any track not reconstructed as a PF candidate with $ d_z < 0.2$ cm and $ d_{xy} < 0.1$ cm	$N_j = 1$ or $H_T < 575$ GeV
Any electron	Single-lepton CR

the M_{T2} variable provides a large rejection of QCD multijet background events. The threshold is increased to 400 GeV for events with large values of H_T to address the larger contribution from this background in this kinematic region. Monojet events are selected with $p_T^{\text{miss}} > 250$ GeV.

- $\Delta\phi_{\min} > 0.3$: the variable $\Delta\phi_{\min}$ is defined as the minimum difference in the azimuthal angle, $\Delta\phi$, between p_T^{miss} and any of the four jets with the largest p_T . Because the $\vec{p}_T^{\text{miss}}$ vector is likely to be aligned to the direction of mismeasured jets, the cut on $\Delta\phi_{\min}$ rejects a large amount of QCD multijet events.
- $|\vec{p}_T^{\text{miss}} - \vec{H}_T^{\text{miss}}|/p_T^{\text{miss}} < 0.5$: while $\vec{p}_T^{\text{miss}}$ is defined as the negative of the vector $\vec{p}_T$ sum of all the reconstructed PF candidates, $\vec{H}_T^{\text{miss}}$ is defined as the negative of the vector $\vec{p}_T$ sum of only the selected jets, as described in Section 4.1.2. This means that the computation of $\vec{H}_T^{\text{miss}}$ excludes jets with $p_T < 30$ GeV and $|\eta| > 2.4$, unlike that of p_T^{miss} . Therefore, this cut protects against a potential bias arising from soft and/or forward jets in the computation of the M_{T2} variable.
- $N_\ell = 0$: as the search is performed in all-hadronic final states, the number of charged leptons, N_ℓ , must be zero. As discussed in Section 4.3, the leptons can either be reconstructed electron or muon candidates; or be PF electron, muon, or hadron candidates. For the latter case, the additional requirement on the transverse mass $m_T(\text{PF candidate}, p_T^{\text{miss}}) < 100$ GeV is imposed. This requirement was found to reduce the background contribution from leptonic W boson decays, while preserving high signal efficiency.

In addition to the above selection criteria, events with anomalous noise in the detector are rejected using dedicated filters [77], as they may bias the computation of p_T^{miss} . Furthermore, because two modules in the HCAL endcap failed during the data taking in 2018, preventing the energy deposited in these modules to be recorded, an additional veto is applied on the corresponding data sets. Events are rejected if they contain, in the detector region defined as $-4.7 < \eta < -1.4$ and $-1.6 < \phi < -0.8$, at least one of the objects listed in Table 4.6.

The signal and background yields passing the event selection are estimated in Sections 4.6 and 4.7, respectively.




Figure 4.3: Illustration of the binning in N_j and N_b for multijet events in the very-low and low H_T regions (*left*), and for the medium, high, and very-high H_T regions (*right*). Each blue square represents a bin in N_j and N_b .

4.5 Event categories

After the selection is applied, events are further categorised into different signal regions (SRs), also referred to as search categories. The aim of this categorisation is to enhance the sensitivity to different signal processes, and to separate the different background sources. For events with at least two jets, the SRs are defined according to H_T , N_j , N_b , and M_{T2} , while for monojet events, they are defined according to N_b and p_T^{jet} .

Events with at least two jets are first categorised into five H_T regions:

1. Very-low H_T : $250 < H_T < 450$ GeV.
2. Low H_T : $450 < H_T < 575$ GeV.
3. Medium H_T : $575 < H_T < 1200$ GeV.
4. High H_T : $1200 < H_T < 1500$ GeV.
5. Very-high H_T : $H_T > 1500$ GeV.

They are then categorised into 14 bins in N_j and N_b for the very-low and low H_T regions, and into 17 bins for the medium, high, and very-high H_T regions, as illustrated in Fig. 4.3. The (H_T, N_j, N_b) regions, hereafter referred to as the *topological regions*, Ω , are further divided into at most 10 bins in M_{T2} . This binning is designed such that, in each topological region, there is at least one expected background event in the highest M_{T2} bin. This requirement defines the lower bound on the last M_{T2} bin. The lower M_{T2} bins are then designed such that they have a width of either 100 or 200 GeV. The resulting 270 signal regions are listed in Tables 4.7 and 4.8.

On the other hand, monojet events are categorised into two bins with either zero or one N_b , which are then categorised into bins in p_T^{jet} . The chosen binning results in 12 SRs, summarised in Table 4.9. Putting together the multijet and monojet categories, the search is therefore binned into 282 mutually exclusive signal regions.

The distribution of signal and background events across the different search categories is discussed in Sections 4.6 and 4.7, respectively.

Table 4.7: Definition of the signal regions for the very-low, low, and medium H_T multijet regions. The labels “j” and “b” are short for N_j and N_b , respectively.

H_T (GeV)	Jet multiplicities	M_{T2} binning (GeV)
[250, 450]	2 – 3j, 0b	[200, 300, 400, ∞]
	2 – 3j, 1b	[200, 300, 400, ∞]
	2 – 3j, 2b	[200, 300, 400, ∞]
	4 – 6j, 0b	[200, 300, 400, ∞]
	4 – 6j, 1b	[200, 300, 400, ∞]
	4 – 6j, 2b	[200, 300, 400, ∞]
	$\geq 7j$, 0b	[200, 300, 500, ∞]
	$\geq 7j$, 1b	[200, 300, ∞]
	$\geq 7j$, 2b	[200, 300, ∞]
	2 – 6j, $\geq 3b$	[200, 300, 400, ∞]
	$\geq 7j$, $\geq 3b$	[200, 300, ∞]
[450, 575]	2 – 3j, 0b	[200, 300, 400, 500, ∞]
	2 – 3j, 1b	[200, 300, 400, 500, ∞]
	2 – 3j, 2b	[200, 300, 400, 500, ∞]
	4 – 6j, 0b	[200, 300, 400, 500, ∞]
	4 – 6j, 1b	[200, 300, 400, 500, ∞]
	4 – 6j, 2b	[200, 300, 400, 500, ∞]
	$\geq 7j$, 0b	[200, 300, 400, ∞]
	$\geq 7j$, 1b	[200, 300, 400, ∞]
	$\geq 7j$, 2b	[200, 300, 400, ∞]
	2 – 6j, $\geq 3b$	[200, 300, 400, 500, ∞]
	$\geq 7j$, $\geq 3b$	[200, 300, 400, ∞]
[575, 1200]	2 – 3j, 0b	[200, 300, 400, 500, 600, 700, 800, 900, 1000, 1100, ∞]
	2 – 3j, 1b	[200, 300, 400, 600, 800, 1000, ∞]
	2 – 3j, 2b	[200, 300, 400, 600, 800, ∞]
	4 – 6j, 0b	[200, 300, 400, 500, 600, 700, 800, 900, 1000, 1100, ∞]
	4 – 6j, 1b	[200, 300, 400, 600, 800, 1000, ∞]
	4 – 6j, 2b	[200, 300, 400, 600, 800, ∞]
	2 – 6j, $\geq 3b$	[200, 300, 400, 600, 800, ∞]
	7 – 9j, 0b	[200, 300, 400, 600, 800, ∞]
	7 – 9j, 1b	[200, 300, 400, 600, 800, ∞]
	7 – 9j, 2b	[200, 300, 400, 600, 800, ∞]
	7 – 9j, 3b	[200, 300, 400, 600, ∞]
	7 – 9j, $\geq 4b$	[200, 300, 400, ∞]
	$\geq 10j$, 0b	[200, 300, 500, ∞]
	$\geq 10j$, 1b	[200, 300, 500, ∞]
	$\geq 10j$, 2b	[200, 300, 500, ∞]
	$\geq 10j$, 3b	[200, 300, ∞]
	$\geq 10j$, $\geq 4b$	[200, ∞]

Table 4.8: Definition of the signal regions for the high and very-high H_T multijet regions. The labels “j” and “b” are short for N_j and N_b , respectively.

H_T (GeV)	Jet multiplicities	M_{T2} binning (GeV)
[1200, 1500]	2 – 3j, 0b	[200, 400, 600, 800, 1000, 1200, ∞]
	2 – 3j, 1b	[200, 400, 600, 800, 1000, 1200, ∞]
	2 – 3j, 2b	[200, 400, 600, 800, 1000, ∞]
	4 – 6j, 0b	[200, 400, 600, 800, 1000, 1200, ∞]
	4 – 6j, 1b	[200, 400, 600, 800, 1000, 1200, ∞]
	4 – 6j, 2b	[200, 400, 600, 800, 1000, ∞]
	2 – 6j, $\geq 3b$	[200, 400, 600, ∞]
	7 – 9j, 0b	[200, 400, 600, 800, 1000, ∞]
	7 – 9j, 1b	[200, 400, 600, 800, 1000, ∞]
	7 – 9j, 2b	[200, 400, 600, 800, ∞]
	7 – 9j, 3b	[200, 400, 600, ∞]
	7 – 9j, $\geq 4b$	[200, 400, ∞]
	$\geq 10j$, 0b	[200, 400, 600, ∞]
	$\geq 10j$, 1b	[200, 400, 600, ∞]
	$\geq 10j$, 2b	[200, 400, 600, ∞]
	$\geq 10j$, 3b	[200, 400, ∞]
	$\geq 10j$, $\geq 4b$	[200, ∞]
[1500, ∞]	2 – 3j, 0b	[400, 600, 800, 1000, 1200, 1400, 1800, ∞]
	2 – 3j, 1b	[400, 600, 800, 1000, 1200, ∞]
	2 – 3j, 2b	[400, ∞]
	4 – 6j, 0b	[400, 600, 800, 1000, 1200, 1400, 1600, ∞]
	4 – 6j, 1b	[400, 600, 800, 1000, 1400, ∞]
	4 – 6j, 2b	[400, 600, 800, ∞]
	2 – 6j, $\geq 3b$	[400, 600, ∞]
	7 – 9j, 0b	[400, 600, 800, 1000, 1400, ∞]
	7 – 9j, 1b	[400, 600, 800, ∞]
	7 – 9j, 2b	[400, 600, 800, ∞]
	7 – 9j, 3b	[400, 800, ∞]
	7 – 9j, $\geq 4b$	[400, ∞]
	$\geq 10j$, 0b	[400, 800, ∞]
	$\geq 10j$, 1b	[400, 800, ∞]
	$\geq 10j$, 2b	[400, ∞]
	$\geq 10j$, 3b	[400, ∞]
	$\geq 10j$, $\geq 4b$	[400, ∞]

 Table 4.9: Definition of the signal regions for the monojet regions. The labels “j” and “b” are short for N_j and N_b , respectively.

Jet multiplicities	p_T^{jet} binning (GeV)
1j, 0b	[250, 350, 450, 575, 700, 1000, 1200, ∞]
1j, $\geq 1b$	[230, 350, 450, 575, 700, ∞]

4.6 Signal signatures

The parameters of several supersymmetric models are constrained in the context of this work. These signal models, described in Section 3.3, target the strong production of sparticles and conserve R -parity. A total of six processes, listed in Table 3.3, are considered. They involve the pair-production of gluinos or squarks that decay into light, beauty, or top flavour quarks together with the LSP, $\tilde{\chi}_1^0$. Moreover, the targeted models have two free parameters: (i) the mass of the directly produced gluinos or squarks, hereafter noted as M , and (ii) the mass of the LSP, $m_{\tilde{\chi}_1^0}$.

In a given signal region, defined in Section 4.5, the expected signal yield, N_{sig} , associated with a given simplified SUSY model with signal hypothesis $(M, m_{\tilde{\chi}_1^0})$ is obtained as

$$N_{\text{sig}}(M, m_{\tilde{\chi}_1^0}) = \sigma_{\text{SUSY}}(M) \times \mathcal{L} \times \epsilon_{\text{sig}}(M, m_{\tilde{\chi}_1^0}), \quad (4.6)$$

where σ_{SUSY} is the cross section of the given simplified model, shown in Fig. 3.3 as a function of the mass M ; $\mathcal{L}$ is the total integrated luminosity accumulated during Run 2 and is equal to 137 fb^{-1} , as discussed in Section 4.2.1; and ϵ_{sig} is the signal efficiency, incorporating the acceptance and selection effects. The signal yield is further corrected with the weights described in Section 4.2.2.

The distribution of the variables M_{T2} , H_T , N_j , and N_b is shown in Fig. 4.4 for simulated backgrounds and representative mass hypotheses, $(M, m_{\tilde{\chi}_1^0})$, for the gluino-mediated signal models. The spectra are shown for events with at least two jets passing the selection described in Section 4.4. For each signal model, two mass hypotheses $(M, m_{\tilde{\chi}_1^0})$ are shown, either with small ($\sim 200 \text{ GeV}$) or large ($\sim 1400 \text{ GeV}$) mass splittings, $|M - m_{\tilde{\chi}_1^0}|$, referred to as the compressed and non-compressed scenarios, respectively. The signal hypotheses are chosen such that they lie within the sensitivity reach of the analysis.

Generally speaking, the sensitivity of the analysis to compressed scenarios is lesser than that to non-compressed scenarios. In the former case, the available kinetic energy for the quarks and LSP is limited because most of the energy of the mother particle is converted into the mass of the LSP. This scenario results in smaller values of M_{T2} and H_T , as can be seen in the upper-left and upper-right plots in Fig. 4.4, respectively. For the non-compressed scenarios, we verify that the signal processes can have high values of M_{T2} . As anticipated in Section 4.1.1, it implies that the sensitivity of the analysis is greater in the tails of the M_{T2} distribution rather than in its bulk.

As discussed in Section 4.5, the definition of the signal regions is designed in such a way that the sensitivity to different signal models is enhanced. For instance, lower H_T categories are more sensitive to compressed mass scenarios, while higher H_T categories are more sensitive to non-compressed mass scenarios, with large values of M . Moreover, the variable N_j is a useful handle to discriminate between signal models involving top quarks and those involving lighter quarks. Indeed, the jet multiplicity is expected to be larger for the former processes, as the top quark decays into bW , with the W boson decaying hadronically. Furthermore, as squark- and gluino-mediated decays give rise to a different number of quarks, it is expected that signal




Figure 4.4: Distributions of M_{T2} (upper left), H_T (upper right), N_j (lower left), and N_b (lower right) for simulated background and signal events, shown for multijet events. Six signal models involving gluino-mediated decays into light-flavour quarks (cyan), beauty quarks (magenta), and top quarks (orange) are shown for non-compressed (solid lines) and compressed (dashed lines) mass scenarios. The normalisation of the signal is increased by a factor 10 for better readability. The plots were produced by the UCSB team.

regions with high N_j will be mostly sensitive to signal processes involving gluinos, while regions with lower N_j will have an increased sensitivity to signal processes involving squarks. Lastly, the sensitivity to models with beauty or top quarks (decaying to bW) is enhanced in categories with large N_b , while the sensitivity to models with light-flavour quarks is higher in categories with small N_b .

Uncertainties

The list of uncertainties in the signal yield is summarised in Table 4.10. In addition to the statistical uncertainties, coming from the limited size of the simulated samples, several sources of systematic uncertainties must be taken into account when extracting the final results.

The largest uncertainties arise from the b-jet tagging efficiency corrections, and from the ISR modelling. The latter is particularly large for signal hypotheses with small mass splittings, $|M - m_{\tilde{\chi}_1^0}|$, for which a boost from ISR is necessary for the final state to be within the detector acceptance. Another significant source of uncertainty comes from the lepton scale factors, which

Table 4.10: Sources of systematic uncertainties in the signal yield, together with their typical magnitude ranges across the search categories.

Source	Value (%)
Limited size of MC samples	1–100
Integrated luminosity	2.3–2.5
ISR modelling	0–30
b-tagging efficiency, heavy flavours	0–40
b-tagging efficiency, light flavours	0–20
Lepton efficiency	0–20
Jet energy corrections	5
μ_R and μ_F	5
Fast simulation p_T^{miss} modelling	0–5

give rise to an uncertainty up to 20% depending on the signal model and the search category. Moreover, the uncertainty in the jet energy corrections are taken into account, as well as the theoretical uncertainties related to the choice of the renormalisation (μ_R) and factorisation (μ_F) scales used in the event generation. Finally, uncertainties in the p_T^{miss} modelling are applied for signal samples using the fast simulation framework, introduced in Section 4.2.2.

All the uncertainties are treated as correlated between the different search regions, except for the statistical uncertainties, which are treated as uncorrelated.

4.7 Background estimation

As anticipated at the beginning of this chapter, the SM processes populating the signal regions can be categorised into three classes: (i) the $Z \rightarrow \nu\bar{\nu}$ background, (ii) the lost-lepton background, and (iii) the QCD multijet background.

As illustrated in Fig. 4.4, the overall background yield is expected to decrease with larger values of M_{T2} and H_T . Moreover, the $Z \rightarrow \nu\bar{\nu}$ background, which consists mainly of Z+jets events, is expected to be the largest source of background in the tail of the M_{T2} distribution and in the high H_T categories. However, the prominence of this background is lesser for events with large jet and b-jet multiplicities. For these events, the W+jets and top processes, which are the two main contributions to the lost-lepton background, dominate. Finally, residual QCD multijet background events that pass the event selection only subsist at high H_T and N_j , which makes this background source the least prominent in this search.

The estimation methods of the $Z \rightarrow \nu\bar{\nu}$, lost-lepton, and QCD multijet backgrounds, as well as their associated uncertainties, are presented in Sections 4.7.1, 4.7.2, and 4.7.3, respectively.

4.7.1 $Z \rightarrow \nu\bar{\nu}$ background

A prominent source of background originates from events in which a Z boson, produced in association with one or more jets, decays invisibly into two neutrinos. Such a process yields the same signatures as the sought-after signals, i.e. final states with a large hadronic activity, with

genuine missing transverse momentum, and without prompt charged leptons. This background, referred to as the $Z \rightarrow \nu\bar{\nu}$ background, is therefore irreducible. Because there isn't inherent production of b quarks, the contribution of these events gets smaller with larger values of N_b .

The number of $Z \rightarrow \nu\bar{\nu}$ background events is estimated from transfer factors applied to the number of $Z \rightarrow \ell^+\ell^-$ ($\ell = e, \mu$) events. For events with at least two jets, the number of predicted $Z \rightarrow \nu\bar{\nu}$ events, $N_{Z \rightarrow \nu\bar{\nu}}^{\text{SR}}$, is obtained in each topological region, $\Omega = (H_T, N_j, N_b)$, and M_{T2} bin as

$$N_{Z \rightarrow \nu\bar{\nu}}^{\text{SR}}(\Omega, M_{T2}) = N_{\text{data}}^{2\ell\text{-CR}}(\Omega) \times \mathcal{P}^{Z \rightarrow \ell^+\ell^-}(\Omega) \times R_{\text{MC}}^{Z \rightarrow \nu\bar{\nu}/Z \rightarrow \ell^+\ell^-}(\Omega) \times k_{Z \rightarrow \nu\bar{\nu}}(M_{T2}), \quad (4.7)$$

and for monojet events, they are obtained in each (p_T^{jet}, N_b) bin as

$$N_{Z \rightarrow \nu\bar{\nu}}^{\text{SR}}(p_T^{\text{jet}}, N_b) = N_{\text{data}}^{2\ell\text{-CR}}(p_T^{\text{jet}}, N_b) \times \mathcal{P}^{Z \rightarrow \ell^+\ell^-}(p_T^{\text{jet}}, N_b) \times R_{\text{MC}}^{Z \rightarrow \nu\bar{\nu}/Z \rightarrow \ell^+\ell^-}(p_T^{\text{jet}}, N_b), \quad (4.8)$$

where

- $N_{\text{data}}^{2\ell\text{-CR}}$ is the number of observed data events in the same-flavour dilepton control sample.
- $\mathcal{P}^{Z \rightarrow \ell^+\ell^-}$ is the purity of the $Z \rightarrow \ell^+\ell^-$ process in the same-flavour dilepton control sample.
- $R_{\text{MC}}^{Z \rightarrow \nu\bar{\nu}/Z \rightarrow \ell^+\ell^-}$ is the transfer factor defined as the ratio of $Z \rightarrow \ell^+\ell^-$ to $Z \rightarrow \nu\bar{\nu}$ events obtained in MC.
- $k_{Z \rightarrow \nu\bar{\nu}}$ is the factor accounting for the distribution of events across the M_{T2} bins, as the previous quantities are obtained in each topological region Ω integrated over M_{T2} . This factor is computed in a hybrid manner, i.e. relying either on data or MC, depending on the available statistics.

These quantities are further discussed in the following.

Dilepton control regions

The number of $Z \rightarrow \ell^+\ell^-$ events is estimated using dedicated dilepton control regions, noted as $2\ell\text{-CR}$. These control regions contain events with exactly two same-flavour (SF) leptons, either ee or $\mu\mu$, that have opposite electric charges. The events are selected using the triggers described in Section 4.2.1. To ensure similar kinematics as the $Z \rightarrow \nu\bar{\nu}$ events, the leading and subleading charged leptons must have $p_T > 100$ GeV and $p_T > 30$ GeV, respectively, and the dilepton system must satisfy $p_T(\ell^+\ell^-) > 200$ GeV. Moreover, the requirement $|m(\ell^+\ell^-) - m(Z)| < 20$ GeV imposes that the dilepton invariant mass, $m(\ell^+\ell^-)$, is in good agreement with the Z boson mass, $m(Z)$, to ensure that most of the selected events arise from the decay of a Z boson. Since the leptons in $Z \rightarrow \ell^+\ell^-$ events must mimic the role of the neutrinos in $Z \rightarrow \nu\bar{\nu}$ events, the lepton $\vec{p}_T$ are vectorially added to the $\vec{p}_T^{\text{miss}}$ vector for the computation of all the kinematic variables. Finally, because leptons are often reconstructed within jets, the closest jet within a cone of size $\Delta R = 0.4$ around each lepton is removed from the event before computing the kinematic quantities.




Figure 4.5: Data to MC simulation comparison of the distribution of the kinematic variables H_T (upper left), N_j (upper right), N_b (lower left), and p_T^{jet} (lower right) for events in the same-flavour dilepton control region. The first three distributions are shown for events with at least two jets and the last one is shown for monojet events. All processes with a top quark are summed and the MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio between the data and MC yields is shown in the lower panel.

The dilepton control regions are binned into the same topological regions Ω as the signal regions, except for the regions with $N_j \geq 7$, which are merged together to maintain reasonable statistical power.

Data to MC simulation comparison of the distribution of the kinematic variables used for the definition of the signal regions are shown in Fig. 4.5, for events in the inclusive dilepton control region. The modelling of the MC is generally good, with some trends nonetheless visible mainly in the H_T and p_T^{jet} spectra. However, as the derivation of the background estimate is primarily data-driven and does not rely directly on these quantities, these effects of the modelling do not affect the final estimate. Figure 4.6 shows the modelling of the M_{T2} variable in the very-low and medium H_T regions. The agreement between data and MC is also generally good.




Figure 4.6: Data to MC simulation comparison of the M_{T2} distribution in the very-low H_T (left) and medium H_T (right) regions, for events in the same-flavour dilepton CR with at least two jets. All processes with a top quark are summed and the MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio between the data and MC yields is shown in the lower panel. The level of agreement between data and MC is similar for the other H_T regions.

The remaining level of mismodelling is taken into account in the systematic uncertainty in the background estimate, discussed below.

Purity

Despite the selection criteria on $p_T(\ell^+\ell^-)$ and $m(\ell^+\ell^-)$, events from other processes than the targeted $Z \rightarrow \ell^+\ell^-$ decays may populate the SF dilepton control regions, 2ℓ -CR. The main source of such events comes from $t\bar{t}$ production, but other processes, such as $t\bar{t}+W$ and $t\bar{t}+Z$, may also contribute. The purity $\mathcal{P}^{Z \rightarrow \ell^+\ell^-}$, defined as the rate of $Z \rightarrow \ell^+\ell^-$ events in the 2ℓ -CR, must therefore be measured to account for this potential contamination.

While the two charged leptons produced in Z boson decays always have the same flavour, the leptons in top processes are produced at the same rate with either the same or a different flavour. This property is exploited to measure the top contamination in the SF 2ℓ -CR, in a fully data-driven manner. The corresponding yield can indeed be obtained from a transfer factor, $R^{\text{SF/DF}}$, applied to the number of events arising from top processes measurement in the different-flavour (DF) dilepton control region. This region is defined with the same criteria as the SF dilepton control region, with the exception that events must contain exactly one electron and one muon, instead of either two electrons or two muons.

The transfer factor, $R^{\text{SF/DF}}$, accounts for the fact that the rate of events with SF and DF leptons, in spite of being equal in theory, may differ experimentally as the triggering, reconstruction, and identification efficiencies for electrons and muons are not exactly the same. The $R^{\text{SF/DF}}$ factor is measured as the ratio of events with SF and DF leptons in a top-enriched region. This region is defined with the same requirements as for the dilepton control regions, but with inverted




Figure 4.7: $R^{\text{SF/DF}}$ factor as a function of the kinematic variables M_{T2} (upper left), H_T (upper right), N_j (lower left), and N_b (lower right). The measured central value $R^{\text{SF/DF}} = 1.06$ is shown as a solid black line, while the statistical and systematic uncertainties are represented with the black and red dashed lines, respectively.

conditions on the dilepton system, i.e. $p_T(\ell^+\ell^-) < 200 \text{ GeV}$ and $|m(\ell^+\ell^-) - m_Z| > 20 \text{ GeV}$. The ratio was measured to be $R^{\text{SF/DF}} = 1.06 \pm 0.03 \text{ (stat)} \pm 0.15 \text{ (syst)}$. It was found to be approximately constant as a function of the kinematic variables used for the event categorisation, as seen in Fig. 4.7, which motivates the choice of keeping this value the same for the background estimation in all search regions. The residual variations are covered by the systematic uncertainty of 15%.

Hybrid M_{T2} shape

For multijet events, the factor $k_{Z \rightarrow \nu\bar{\nu}}$ distributes the predicted background events across the M_{T2} bins in a given topological region Ω . It is obtained from the M_{T2} shape taken either from data or MC simulation, depending on the available statistics in the given bin. The introduction of this hybrid method minimises the use of MC, which may be subjected to mismodelling, and




Figure 4.8: Sketch illustrating the M_{T2} distribution in the same-flavour dilepton control region in a given topological region Ω . The main elements of the hybrid method used to extrapolate the M_{T2} shape to the $Z \rightarrow \nu\bar{\nu}$ estimate are reported.

makes the estimation of the $Z \rightarrow \nu\bar{\nu}$ background almost completely data-driven.

The *hybrid threshold* refers to the M_{T2} bin boundary below (above) which data (MC) is used to extrapolate the M_{T2} shape, as illustrated in Fig. 4.8. It is determined, in each topological region Ω , using simulated events in the same-flavour dilepton control region normalised to the total integrated luminosity. Starting from the highest M_{T2} bin, the lower bins are merged until the accumulated number of events in the merged bins is equal to at least 50. The hybrid threshold corresponds to the lower boundary of the lowest M_{T2} bin that is merged. Then,

- For bins below the hybrid threshold, typically corresponding to the bulk of the M_{T2} distribution, the M_{T2} shape is built using $Z \rightarrow \ell^+\ell^-$ data, corrected by the $R_{\text{MC}}^{Z \rightarrow \nu\bar{\nu}/Z \rightarrow \ell^+\ell^-}$ factor.
- For bins above the hybrid threshold, typically corresponding to the tail of the M_{T2} distribution, the M_{T2} shape is distributed using $Z \rightarrow \nu\bar{\nu}$ MC events, after these events are normalised to the total data yield across the merged bins.

In case the number of events in the highest M_{T2} bin is larger than 50, the hybrid threshold is taken to be infinity, and only data are used to extrapolate the M_{T2} shape.

In each H_T region, the M_{T2} shape for simulated $Z \rightarrow \ell^+\ell^-$ and $Z \rightarrow \nu\bar{\nu}$ events is found to be independent of the value of N_b . Representative examples are shown in Fig. 4.9 for events in the low H_T region with $2 \leq N_j \leq 3$. A similar behaviour was observed with respect to N_j for events in the very-high H_T region. As a consequence, the M_{T2} shape is obtained integrating over N_b and N_j for events in the very-high H_T region, and integrating over N_b for events in the other H_T regions.

The hybrid treatment is validated by comparing the M_{T2} shape of simulated $Z \rightarrow \nu\bar{\nu}$ events to that of data $Z \rightarrow \ell^+\ell^-$ events from the dilepton control region. Examples of the closure are shown in Fig. 4.10 in two representative topological regions. The good agreement of the




Figure 4.9: Shape of the M_{T2} distribution for $Z \rightarrow \nu\bar{\nu}$ (left) and $Z \rightarrow \ell^+\ell^-$ (right) simulated events for different N_b regions. The comparison is shown for a representative region with $450 < H_T < 575$ GeV and $2 \leq N_j \leq 3$. The lower panel shows the ratio of the $N_b = 1$ and $N_b = 2$ entries to the $N_b = 0$ entries.




Figure 4.10: Comparison of the M_{T2} shape from $Z \rightarrow \ell^+\ell^-$ data and $Z \rightarrow \nu\bar{\nu}$ MC for events in the very-low H_T (left) and medium H_T (right) regions with $2 \leq N_j \leq 3$ and $N_b = 0$. The errors are statistical.

shapes validates the hybrid approach. The residual discrepancies are covered by the systematic uncertainties, described in the following.

Uncertainties

Different sources of uncertainties in the $Z \rightarrow \nu\bar{\nu}$ background estimate must be taken into account when deriving the final results. They are summarised in Table 4.11.

The statistical uncertainties, resulting from the limited size of the data and MC samples, is obtained as the Poisson error on the number of events in the given search region. They have a

Table 4.11: Sources of systematic uncertainties in the $Z \rightarrow \nu\bar{\nu}$ background estimate, together with their typical magnitude ranges across the search categories.

Source	Value (%)
Limited size of data control samples	5–100
Limited size of MC samples	0–50
Lepton efficiency	0–5
Jet energy corrections	0–5
μ_R and μ_F variation	0–5
$R^{\text{SF/DF}}$ (stat)	0–3
$R^{\text{SF/DF}}$ (syst)	0–15
M_{T2} shape	0–40

large contribution to the total uncertainty across the search categories.

The systematic uncertainties in the transfer factor, $R_{\text{MC}}^{Z \rightarrow \nu\bar{\nu}/Z \rightarrow \ell^+\ell^-}$, the purity, $\mathcal{P}^{Z \rightarrow \ell^+\ell^-}$, and the M_{T2} shape, $k_{Z \rightarrow \nu\bar{\nu}}$ are described below.

First of all, the uncertainty in the lepton scale factors are the largest uncertainty in the transfer factor, $R_{\text{MC}}^{Z \rightarrow \nu\bar{\nu}/Z \rightarrow \ell^+\ell^-}$. These corrections and their uncertainties are provided by the responsible working group at CMS. The associated uncertainties are propagated to the final results and result in a systematic uncertainty up to 5%, depending on the signal region. Second of all, the uncertainty in the jet energy corrections, introduced in Section 4.3, also contribute to the total uncertainty in the transfer factor. Finally, theoretical uncertainties related to the choice of the renormalisation (μ_R) and factorisation (μ_F) scales used in the event generation are also taken into account.

The main uncertainty in the purity, $\mathcal{P}^{Z \rightarrow \ell^+\ell^-}$, is associated with the factor $R^{\text{SF/DF}}$, discussed in Section 4.7.1. The statistical and systematic uncertainties in this factor are 3% and 15%, respectively.

Finally, for M_{T2} bins above the hybrid threshold, an additional systematic uncertainty associated with the factor $k_{Z \rightarrow \nu\bar{\nu}}$ is applied. Both theoretical and instrumental effects are accounted for, in particular the modelling of the M_{T2} variable in simulation, causing this shape uncertainty reach values as large as 40% in the highest M_{T2} bins.

All the uncertainties are treated as uncorrelated between the different search regions, apart from the uncertainties in $R^{\text{SF/DF}}$, which are treated as correlated.

4.7.2 Lost-lepton background

The second largest source of background, referred to as the lost-lepton background, originates from processes in which a prompt charged lepton (e or μ) is produced from the decay of a W boson. While this source of background is largely suppressed by imposing the lepton veto described in Section 4.4, such background events may contaminate the signal regions in case the lepton is lost. A lepton is defined as lost if it is out of acceptance, not reconstructed, not identified, or not isolated. Because leptonic W boson decays also give rise to genuine $p_{\text{T}}^{\text{miss}}$ through the neutrinos, these events can fake all-hadronic final states with large $p_{\text{T}}^{\text{miss}}$. This

background arises mostly from W +jets and $t\bar{t}$ +jets processes, with smaller contributions from single top, $t\bar{t}+W$, and $t\bar{t}+Z$ processes. Because of the decay of the top quark into bW , this background source is more prominent in event categories with a large b -jet multiplicity.

In each search region, the number of events, N_{LL}^{SR} , associated with the lost-lepton background is estimated by means of a transfer factor applied to the yield in the corresponding single-lepton control region, 1ℓ -CR, defined below. For multijet events, the background yield in each topological region, Ω , and in each M_{T2} bin is computed as

$$N_{LL}^{SR}(\Omega, M_{T2}) = N_{data}^{1\ell-CR}(\Omega, M_{T2}) \times R_{MC}^{0\ell/1\ell}(\Omega, M_{T2}) \times k_{LL}(M_{T2}), \quad (4.9)$$

while for monojet events, it is computed in each (p_T^{jet}, N_b) bin as

$$N_{LL}^{SR}(p_T^{\text{jet}}, N_b) = N_{data}^{1\ell-CR}(p_T^{\text{jet}}, N_b) \times R_{MC}^{0\ell/1\ell}(p_T^{\text{jet}}, N_b), \quad (4.10)$$

where

- $N_{data}^{1\ell-CR}$ is the number of data events in the single-lepton control region.
- $R_{MC}^{0\ell/1\ell}$ is the transfer factor, obtained from simulation as the ratio of the number of events in the signal region (0ℓ) to that in the single-lepton control region (1ℓ). Because the contribution from all processes contributing to the lost-lepton background are summed, a single transfer factor for all is computed for a given search region.
- k_{LL} is the factor used to distribute the events across the M_{T2} bins.

These quantities are further described in the following.

Single-lepton control regions

The single-lepton control regions, 1ℓ -CR, are defined using the same trigger and selection criteria as for the signal regions, except for the lepton veto, which is inverted by imposing the presence of exactly one lepton (e or μ) in the final state. To reduce the potential contamination from signal events in these regions, the transverse mass, computed with the lepton candidate and $\vec{p}_T^{\text{miss}}$, must obey $m_T(\text{lepton}, \vec{p}_T^{\text{miss}}) < 100$ GeV. For monojet events, the single-lepton control regions are defined using the same categorisation in p_T^{jet} and N_b as the signal regions, defined in Section 4.5; while for multijet events, the same categorisation in H_T , N_j , N_b , and M_{T2} as the signal regions is used, except for regions with large jet multiplicity ($N_j \geq 7$), which are merged to preserve statistical precision. The events in the 1ℓ -CR are up to twice as many as the events in the corresponding signal region.

Figure 4.11 shows the shape of the distribution of the category variables H_T , N_j , N_b , and p_T^{jet} obtained from data and MC simulation, for events in inclusive single-lepton control regions with either one or more jets. The modelling of these kinematic variables in simulation is generally good, although some trends are visible, especially in the H_T and p_T^{jet} spectra. However, as for the $Z \rightarrow \nu\bar{\nu}$ background estimation, these trends do not affect the estimation of the background yield, because the background estimate is computed primarily in a data-driven manner, as




Figure 4.11: Comparison of data and MC distributions of H_T (upper left), N_j (upper right), and N_b (lower left), for events with at least two jets, and of p_T^{jet} (lower right), for events with exactly one jet. The distributions are shown for events in the single-lepton control region. All processes with a top quark are summed. The MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio of data to MC distributions is shown in the lower panel.

further discussed below, and because these kinematic variables are solely used to define the search regions.

The modelling of the M_{T2} variable in MC events is validated against data in the different H_T regions, as shown in Fig. 4.12 for events in the very-low and medium H_T ranges. A good agreement between the shape in data and simulation is obtained in all regions. The residual discrepancy is accounted for with a dedicated systematic uncertainty, described below.

Finally, the modelling of the W+jets and top processes is validated in regions enriched in these processes via the selection $N_b = 0$ and $N_b \geq 2$, respectively. The modelling of the kinematic variables is similar for the separate processes, as illustrated with the data to MC comparison of the H_T distribution in Fig. 4.13.




Figure 4.12: Comparison of data and MC distributions of the M_{T2} variable for events in the single-lepton control region in the very-low H_T (left) and medium H_T (right) regions. All processes with a top quark are summed. The MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio of data to MC distributions is shown in the lower panel.




Figure 4.13: Comparison of data and MC distributions of the H_T variable for multijet events in the single-lepton control region with $N_b = 0$ (left) and $N_b \geq 2$ (right). All processes with a top quark are summed. The MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio of data to MC distributions is shown in the lower panel.

Hybrid M_{T2} shape

In a similar manner as for the $Z \rightarrow \nu\bar{\nu}$ background, the M_{T2} shape of the lost-lepton background estimate is obtained in a hybrid manner, relying either fully on data or on both data and MC, depending on the available statistics. For multijet events in the single-lepton control region and topological region Ω , the highest M_{T2} bin is merged with the lower bins until the sum of the expected events in MC over the merged bins is at least equal to 50. The lower boundary of the

lowest M_{T2} bin that is merged corresponds to the hybrid threshold, illustrated in Fig. 4.8 in the context of the $Z \rightarrow \nu\bar{\nu}$ estimate. Then,

- For M_{T2} bins below the hybrid threshold, the quantities $N_{\text{data}}^{1\ell\text{-CR}}$ and $R_{\text{MC}}^{0\ell/1\ell}$ are obtained directly in specific M_{T2} bin, and the factor k_{LL} is equal to one.
- For M_{T2} bins above the hybrid threshold, the quantities $N_{\text{data}}^{1\ell\text{-CR}}$ and $R_{\text{MC}}^{0\ell/1\ell}$ are obtained from events in the control region integrated over the merged bins, and the factor k_{LL} is used to restore the distribution of the events across the M_{T2} bins. This factor is computed in MC simulation as the ratio of the signal region yield in a given M_{T2} bin to that integrated over the merged bins.

Uncertainties

Both statistical and systematic uncertainties in the background estimate must be assessed and propagated to the final results. The corresponding list of uncertainties is provided in Table 4.12. The systematic uncertainties in the transfer factor, $R_{\text{MC}}^{0\ell/1\ell}$, and the M_{T2} shape factor, k_{LL} , are described in the following.

The systematic uncertainties associated with the transfer factor, $R_{\text{MC}}^{0\ell/1\ell}$, primarily arise from the corrections applied to MC events to improve the modelling of the various kinematic and topological variables in simulation. Corrections are applied on the jet energy, the efficiency of the lepton identification, and the efficiency of b-jet tagging. The responsible working groups at CMS provide uncertainties in these corrections, which are propagated to the final background estimate, reaching values as high as 10%, depending on the search category. Moreover, the difference in efficiency between data and MC of the cut on m_T , used to define the 1 ℓ -CR, is accounted for. Theoretical uncertainties related to the choice of the renormalisation (μ_R) and factorisation (μ_F) scales used in the event generation are also taken into account. Finally, it was found that $t\bar{t}b\bar{b}$ processes were underpredicted in simulation with respect to $t\bar{t}j\bar{j}$ processes [78]. A correction factor of 1.7 ± 0.5 is therefore applied as an additional weight on $t\bar{t}b\bar{b}$ events. The uncertainty in this factor is propagated to the final estimate and reaches values as high as 25%, depending on the search category.

Another source of systematic uncertainty is associated with the M_{T2} shape extrapolation via the factor k_{LL} . This uncertainty applies to the search categories in which M_{T2} shape is taken from MC ($k_{\text{LL}} \neq 1$). It comprises the MC uncertainties as well as the effect of the M_{T2} mismodelling discussed above and can be as high as 40%.

The shape and statistical uncertainties are treated as uncorrelated between the different search regions, while the other uncertainties are treated as correlated.

Table 4.12: Sources of systematic uncertainties in the lost-lepton background estimate, together with their typical magnitude ranges across the search categories.

Source	Value (%)
Limited size of data control samples	5–100
Limited size of MC samples	0–50
e/μ efficiency	0–10
τ efficiency	0–3
b-tagging efficiency	0–3
Jet energy corrections	0–5
m_T (lepton, $\vec{p}_T^{\text{miss}}$) selection efficiency	0–3
μ_R and μ_F variation	0–5
$t\bar{t}b\bar{b}/t\bar{t}j\bar{j}$ weight	0–25
M_{T2} shape (if $k_{LL} \neq 1$)	0–40

4.7.3 QCD multijet background

The last background populating the signal regions originates from QCD multijet events without genuine missing transverse momentum. This background is mostly instrumental and is a consequence of jet momentum mismeasurements or detector noise, which fake the signatures of missing transverse momentum. As discussed in Sections 4.1.1 and 4.4, this background source is drastically reduced by the cuts on M_{T2} and $\Delta\phi_{\min}$. Its contribution to the total SM background is thus minor compared to the $Z \rightarrow \nu\bar{\nu}$ and lost-lepton backgrounds, as can be seen in the distributions shown in Fig. 4.4. The event categories in which this background is the most prominent are those with high H_T and jet multiplicity.

The estimation of this background and its uncertainty was fully taken care of by the UCSB and UCSD teams, and is therefore not part of this thesis. Further details on the method employed to compute the associated yield can be found in Ref. [68].

4.8 Results and interpretation

In each of the 282 search regions, listed in Section 4.5, the number of observed data events is compared to the sum of the yields from the $Z \rightarrow \nu\bar{\nu}$, lost-lepton, and QCD multijet backgrounds. As discussed in Section 4.1.1, an excess of data over the SM expectation in the tail of the M_{T2} distribution would be a sign for the presence of new physics in the data. Moreover, depending on the production and decay mode of the potential supersymmetric particles, it is expected that the excess of events, if any, is observed simultaneously in different search categories.

Figures 4.14 – 4.16 show the comparison of the observed data events with the total background prediction, obtained as the sum of the three background estimates, for an integrated luminosity equal to 137 fb^{-1} accumulated during the Run 2 of data taking. The different figures show the results in the monojet and the different H_T regions. In each figure, the yields obtained from data and from the three background components are reported in each of the signal categories. No significant excess of data over the SM prediction is observed in any of the search regions.




Figure 4.15: Comparison of the estimated background to the observed data events in each topological region Ω and M_{T2} bin in the low H_T (upper) and medium H_T (lower) regions. The hatched bands represent the full uncertainty in the background prediction. The labels “j” and “b” are short for N_j and N_b , respectively, and the M_{T2} bins given on the x -axis are given in units of GeV. The lower panel shows the ratio of data to the background prediction. The multijet estimates (yellow) were provided by UCSB.




Figure 4.16: Comparison of the estimated background to the observed data events in each topological region Ω and M_{T2} bin in the high H_T (upper) and very-high H_T (lower) regions. The hatched bands represent the full uncertainty in the background prediction. The labels “j” and “b” are short for N_j and N_b , respectively, and the M_{T2} bins given on the x -axis are given in units of GeV. The lower panel shows the ratio of data to the background prediction. The multijet estimates (yellow) were provided by UCSB.




Figure 4.17: Pull distribution obtained from the comparison of the observed data with the background prediction in each of the 282 search regions. The red curve shows the result from a fit with a gaussian function with mean $\mu = 0.04$ and standard deviation $\sigma = 0.8$.

The compatibility of the observed data with the estimated SM background is assessed with the pull distribution. In each search region, the pull is defined as

$$\text{pull} = \frac{n_{\text{data}} - n_{\text{background}}}{\sqrt{\sigma_{\text{data}}^2 + \sigma_{\text{background}}^2}}, \quad (4.11)$$

where n_{data} and $n_{\text{background}}$ represent the observed data and estimated background yields in the given search region, respectively, and where σ_{data} and $\sigma_{\text{background}}$ are the associated total uncertainties. The resulting pull distribution is shown in Fig. 4.17. It is found to be compatible with a gaussian function with mean close to 0 and standard deviation close to 1. It thus confirms the compatibility of the observed data with the SM prediction, and hence the absence of a significant sign for the presence of new physics in the data.

The results presented in Figs. 4.14 – 4.16 are obtained using the $Z \rightarrow \nu\bar{\nu}$, lost-lepton, and QCD multijet background prediction and their uncertainties as obtained with the respective estimation methods, described in Section 4.7, and without a fit to the observed data. Possible correlations between the different signal regions are thus not taken into account. In the following, fits to the data are performed in order to constrain the parameter space of signal models by means of upper exclusion limits.

The new physics models constrained in this work are the simplified supersymmetric models presented in Table 3.3. The computation of the signal yields and their uncertainties are described in Section 4.6. A total of six processes targeting the decays of gluinos or squarks into light, beauty, or top flavour quarks, and the LSP, $\tilde{\chi}_1^0$, are considered. The parameter space is defined with the mass, M , of the directly produced gluinos or squarks and with the mass, $m_{\tilde{\chi}_1^0}$, of the LSP. As discussed in Section 4.6, the distribution of the expected signal yield across the search categories depends on the signal model and the mass hypothesis, noted as $(M, m_{\tilde{\chi}_1^0})$.

The exclusion limits are computed following the procedure described in Ref. [79]. For each

signal model and each signal mass hypothesis, a maximum likelihood function, $\mathcal{L}(r, \vec{\theta})$, is constructed as the product of Poisson probabilities of the 282 signal regions and of the 3 years of data taking, using the signal strength, r , as the parameter of interest and the full set of nuisance parameters, $\vec{\theta}$. The nuisance parameters account for the different sources of uncertainty, described in Sections 4.6 and 4.7 for the signal and background components, respectively. The exclusion limits are computed using the modified frequentist CL_s approach [80]. The CL_s is defined as

$$\text{CL}_s = \frac{\text{CL}_{s+b}}{\text{CL}_b}, \quad (4.12)$$

where CL_{s+b} and CL_b are the p -values computed for the alternative, H_1 , signal+background hypothesis and the null, H_0 , background-only hypothesis, respectively:

$$\text{CL}_{s+b} = P(q_r \geq q_r^{\text{obs}} | H_1), \quad (4.13)$$

and

$$\text{CL}_b = P(q_r \geq q_r^{\text{obs}} | H_0), \quad (4.14)$$

where P denotes the probability, q_r is the profile likelihood ratio [79] used as the test statistic, and q_r^{obs} is the observed value of q_r for each of the tested values of r . In the context of this work, the probability density functions are obtained using the asymptotic approximation [79]. A given signal hypothesis is considered to be excluded at 95% confidence level (CL) if $\text{CL}_s < 5\%$.

Figures 4.18 and 4.19 show the upper exclusion limits at 95% CL on the cross section for the supersymmetric models involving gluino-mediated and squark-mediated decays, respectively. The black and red curves show the observed and expected upper exclusion limits on the signal strength, r . The area enclosed by the curves represents the excluded region. The coloured z -axis indicates the observed upper exclusion limits on the cross section.

The presence of products from gluino-mediated decays in the data is excluded at 95% CL for gluino masses up to approximately $M = 2200$ GeV and LSP masses up to approximately $m_{\tilde{\chi}_1^0} = 1500$ GeV, as can be seen in Fig. 4.18. For squark-mediated decays, the corresponding exclusion limits are up to approximately $M = 1700$ GeV and $m_{\tilde{\chi}_1^0} = 900$ GeV, as shown in Fig. 4.19. The sensitivity of the analysis to signal hypotheses with larger values of M is primarily limited by that the new physics cross sections, shown in Fig. 3.3, are rapidly falling with larger values of M . On the other hand, the sensitivity to scenarios with larger values of $m_{\tilde{\chi}_1^0}$ is limited because of kinematic reasons. Indeed, for a fixed value of M , the larger is the value of $m_{\tilde{\chi}_1^0}$, the smaller is the value of H_T . Therefore, the sensitivity to signals with large $m_{\tilde{\chi}_1^0}$ is typically driven by the medium H_T categories, defined in Section 4.5, as opposed to the signals with small $m_{\tilde{\chi}_1^0}$, whose sensitivity is driven by the very-high H_T categories. Because the background yield is larger for smaller H_T categories, the sensitivity of the analysis decreases with larger values of $m_{\tilde{\chi}_1^0}$. This argument also explains why the limits on the cross sections are found to be more stringent for scenarios with large mass splitting, $|M - m_{\tilde{\chi}_1^0}|$, than for those with small splittings.

The difference in the sensitivity between the gluino-mediated and squark-mediated supersymmetric models is a consequence of the different associated jet multiplicities, N_j , which tend to be larger in the former models than in the latter. Because the background yield decreases




Figure 4.18: The 95% CL upper exclusion limits for the gluino-mediated models, where the gluinos decay into light quarks (*upper left*), beauty quarks (*upper right*), or top quarks (*lower*), shown as a function of the masses of the gluino and the LSP. The thick black curve corresponds to the observed limit, while the thin black ones show the theoretical uncertainties in the signal cross section. The thick red curve corresponds to the median expected limits, while the thin red ones show the corresponding $\pm 1\sigma$ and $\pm 2\sigma$ ranges, where σ is the standard deviation. The coloured axis shows the observed upper limit on the signal cross section. The Feynman diagrams of the corresponding models are shown in Fig. 3.4.




Figure 4.19: The 95% CL upper exclusion limits for the squark-mediated models with pair-production of light-flavour squarks (*upper left*), beauty squarks (*upper right*), or top squarks (*lower*), shown as a function of the masses of the corresponding squark and the LSP. The thick black curve corresponds to the observed limit, while the thin black ones show the theoretical uncertainties in the signal cross section. The thick red curve corresponds to the median expected limits, while the thin red ones show the corresponding $\pm 1\sigma$ and $\pm 2\sigma$ ranges, where σ is the standard deviation. The coloured axis shows the observed upper limit on the signal cross section. In the lower figure, the grey dashed line shows the scenario in which the mass of the top squark equals the sum of the mass of the top quark and the mass of the LSP. The Feynman diagrams of the corresponding models are shown in Fig. 3.5.

Table 4.13: Highest 95% CL exclusion limits on directly produced sparticle mass, M , and on the LSP mass, $m_{\tilde{\chi}_1^0}$, obtained for the different simplified SUSY models. The limits are presented from this analysis (using 137 fb^{-1}) and from the previous iteration of this analysis [66] (using 36 fb^{-1}). The relative difference between the two sets of results is also reported.

Simplified model	Highest limit on M (GeV)			Highest limit on $m_{\tilde{\chi}_1^0}$ (GeV)		
	137 fb^{-1}	36 fb^{-1}	Relative diff.	137 fb^{-1}	36 fb^{-1}	Relative diff.
$pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow q\bar{q}\tilde{\chi}_1^0$	1970	1860	6%	1200	1100	9%
$pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow b\bar{b}\tilde{\chi}_1^0$	2250	2025	11%	1525	1400	9%
$pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$	2250	1900	18%	1250	1010	24%
$pp \rightarrow \tilde{q}\tilde{q}, \tilde{q} \rightarrow q\bar{q}\tilde{\chi}_1^0$	1710	1550	10%	870	775	12%
$pp \rightarrow \tilde{b}\tilde{b}, \tilde{b} \rightarrow b\bar{b}\tilde{\chi}_1^0$	1240	1175	5%	700	590	19%
$pp \rightarrow \tilde{t}\tilde{t}, \tilde{t} \rightarrow t\bar{t}\tilde{\chi}_1^0$	1200	1070	12%	580	550	5%

with increasing number of jets, the search is more sensitive to gluino-mediated decays. Furthermore, the limits on models involving top or beauty quarks are more stringent than those on models involving light-flavour quarks. This behaviour results from that the categories with a large b-jet multiplicity, N_b , drive the sensitivity in the former models, which gives an overall better sensitivity since the background yield is smaller in these categories.

The highest limits obtained on the masses M and $m_{\tilde{\chi}_1^0}$ are reported for each signal model in Table 4.13 for this analysis, using 137 fb^{-1} of data, as well as for the previous iteration of this analysis [66], using 35 fb^{-1} of data. The limits are improved by up to 24% in this search and are among the most stringent to date for all the models.

Other searches conducted at the ATLAS and CMS experiments were performed using the same luminosity and the same final state signatures as the present analysis. A detailed overview can be found in Ref. [3]. An example is shown in Fig. 4.20, where limits from different searches targeting models of pair-production of beauty squarks are presented for both the CMS and ATLAS Collaborations. For instance, a search performed by the CMS Collaboration employed a similar strategy as the one used in the present analysis, but leveraged on the variable H_T^{miss} in place of the M_{T2} variable [81]. On the side of the ATLAS Collaboration, a search was performed using different techniques to enhance the sensitivity in signatures involving b-jets, using multivariate-based techniques as well as a dedicated displaced-vertex-finding algorithm [82]. These two searches have a similar sensitivity as the one from this work.

Because the large effort dedicated to searches for promptly decaying sparticles has not led to a discovery of supersymmetry yet, searches for long-lived sparticles have been gaining more and more attention, as they can give access to so far unprobed regions of the parameter space [85]. As an example, the framework of the M_{T2} search was used to search for long-lived charginos decaying into a neutralino within the tracker volume [68]. Because the charged particle decays into an invisible particle, the associated track is not reconstructed up the edge of the tracker. This signature, referred to as disappearing track, is exploited in the design of the search. A comparison of the sensitivity of this analysis with the regular M_{T2} analysis is however not straightforward as they constrain different signal models.




Figure 4.20: Comparison of the limits from different searches targeting SUSY processes involving the pair-production of beauty squarks. The limits are shown from several searches performed at the CMS (left) and ATLAS (right) Collaborations. The corresponding references are provided in the legend of the figures. In particular, the limits from the present analysis and from the CMS analysis based on H_T^{miss} [81] are shown in red and blue in the left figure, respectively. The limits from the ATLAS search [82] targeting the same process as in the present analysis are shown in dark blue in the right figure. The left figure is taken from Ref. [83] and the right one from Ref. [84].

4.9 Summary

A search for supersymmetry in final states with large hadronic activity and transverse momentum imbalance, and without leptons has been presented. The search exploits proton-proton collision data collected at $\sqrt{s} = 13$ TeV corresponding to an integrated luminosity of 137 fb $^{-1}$. The M_{T2} variable is used as the discovery variable, as it is an efficient discriminant between signal and background processes, taking advantage of their different event topologies. The events are categorised in about 300 signal regions, to maximise the sensitivity of the analysis to different signal models. The main background processes, originating from Z+jets, W+jets, and $t\bar{t}$ +jets processes are estimated using dedicated data-driven techniques. No significant excess of observed data over the SM background prediction is observed in any of the signal regions. Upper exclusion limits are derived at 95% CL for various simplified supersymmetric models. These limits are improved up to 24% with respect to the previous iteration of the analysis, and are among the most stringent to date for all the models. The results from this work are published in the *European Physical Journal C* [68].

Part III

Search for new physics in the neutrino sector

Chapter 5

The physics of heavy neutrinos

Heavy neutrinos, also referred to as heavy neutral leptons (HNLs), are hypothetical particles predicted to be the right-handed chiral partners of the SM neutrinos. The existence of these particles could provide answers to some of the fundamental questions left unanswered in the SM. Therefore, heavy neutrinos represent a possible pathway for discovery that is actively investigated. In this chapter, we present the key elements to understand the physics of heavy neutrinos and their potential implications. We take as a starting point the observation of a broad range of neutrino flavour oscillation phenomena. As argued in Section 5.1, these observations imply that at least two of the three neutrino masses must be non-zero, in contradiction to the zero masses predicted by the SM. In Section 5.2, we discuss how heavy neutrinos are involved in a possible mechanism, known as the Type I seesaw mechanism, to make the SM neutrinos massive. In this section, the key properties of these particles are also summarised. We then present in Section 5.3 the neutrino minimal standard model (ν MSM), which incorporates the heavy neutrinos, and we describe its cosmological consequences. A step closer to the experiment is taken in Section 5.4, which describes the phenomenology underlying the search for heavy neutrinos presented in Chapter 6. Finally, an overview of the existing constraints on the parameter space is given in Section 5.5.

5.1 Active neutrino oscillations

In Section 1.2, we elaborated on the fact that neutrinos are assumed to be massless in the standard model. However, as discussed below, this expectation is refuted by the experimental observation that the neutrinos can oscillate in flavour, since such phenomenon is only possible if neutrinos are massive.

The oscillation phenomenon refers to the process through which neutrinos experience a flavour transition as they propagate over large distances. It is a consequence of that the neutrino physical eigenstates that propagate are not equivalent to the neutrino flavour eigenstates that enter the lepton doublets in Eq. (1.2). In fact, the physical eigenstates are expressed as a linear combination of the flavour eigenstates. As a result, a neutrino produced with flavour α may interact to produce a charged lepton with a different flavour β . The first evidence for neutrino

oscillations was obtained from studies of solar [86] and atmospheric [87] neutrinos, and could explain the deficit in the number of observed ν_e and ν_μ , respectively.

It can be shown that the oscillation probability is a function of the difference of the squared neutrino masses [1]. To explain simultaneously the solar and atmospheric neutrino deficits, there must therefore be at least two massive neutrinos. This is in direct contradiction with the prediction of zero masses in the SM. Although the absolute mass scale of the neutrinos has not yet been measured, cosmological observations constrain them to be orders of magnitude lighter than the other fundamental fermions, since the sum of the neutrino masses is found to be less than 0.12 eV [88].

The fact that neutrinos have very small, yet non-zero, masses is one of the most striking examples of evidence for physics beyond the SM. Indeed, as argued in Section 1.2, the masses of the neutrinos cannot be accommodated in the standard model, as the Yukawa interactions involve both left-handed and right-handed chiral states, while the latter have never been observed experimentally for neutrinos. A mechanism must therefore be derived and validated to explain both the existence and smallness of the neutrino masses. Possible mechanisms that do not rely on the addition of new particles into the theory are discussed in Ref. [89]. In the next section, instead, we describe a mechanism that involves the introduction of new particles, the heavy neutrinos.

5.2 Heavy neutrinos as a solution to the neutrino mass puzzle

The approach taken in this work to provide a mechanism for the active neutrino mass generation consists in extending the particle content of the model with n right-handed neutrinos, ν_R^i , with $i = 1, \dots, n$ the flavour index [89]. These neutrinos are singlets under the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetry and are referred to as *sterile* particles. In the following, we describe the process, referred to as the Type I seesaw mechanism [90], from which the prediction of the existence of heavy neutrinos follows. Other models introducing heavy neutrinos are discussed for instance in Refs. [91, 92].

Masses of the right-handed neutrinos can be introduced to the Lagrangian via two possible terms:

1. Dirac mass term: it originates from the spontaneous breaking of the electroweak symmetry, similarly to the other fermions. When suppressing the flavour indices, this mass term has the form

$$\mathcal{L}_{\text{Dirac}} = -m_D \bar{\nu}_L \nu_R + \text{h.c.} = -\frac{1}{2} (m_D \bar{\nu}_L \nu_R + m_D^T \bar{\nu}_R^c \nu_L^c) + \text{h.c.}, \quad (5.1)$$

where m_D is the matrix of Dirac masses and where we used the fact that the charge conjugation operator acts equivalently as the parity operator, i.e. $\bar{\nu}_L = \bar{\nu}_R^c$ and $\nu_R = \nu_L^c$.

2. Majorana mass term: because ν_R^i are singlets under the SM gauge transformations, additional terms formed uniquely out of these fields can be added to the Lagrangian without

breaking the gauge invariance of the theory. The second mass term is expressed as

$$\mathcal{L}_{\text{Majorana}} = -M_{\text{M}} \bar{\nu}_{\text{R}}^{\text{c}} \nu_{\text{R}} + \text{h.c.}, \quad (5.2)$$

where M_{M} is the matrix of Majorana masses. Such a term couples a particle with the corresponding antiparticle, which is only possible if the particle is neutral. Moreover, since it involves two neutrino fields, it violates lepton number conservation by two units.

The most general renormalisable Lagrangian for neutrino masses includes both the Dirac and Majorana terms:

$$\mathcal{L}_{\text{mass}} = -\frac{1}{2} (m_{\text{D}} \bar{\nu}_{\text{L}} \nu_{\text{R}} + m_{\text{D}}^{\text{T}} \bar{\nu}_{\text{R}}^{\text{c}} \nu_{\text{L}}^{\text{c}} + M_{\text{M}} \bar{\nu}_{\text{R}}^{\text{c}} \nu_{\text{R}}) + \text{h.c.}, \quad (5.3)$$

which can be rewritten as

$$\mathcal{L}_{\text{mass}} = -\frac{1}{2} \begin{pmatrix} \bar{\nu}_{\text{L}} & \bar{\nu}_{\text{R}}^{\text{c}} \end{pmatrix} \begin{pmatrix} 0 & m_{\text{D}} \\ m_{\text{D}}^{\text{T}} & M_{\text{M}} \end{pmatrix} \begin{pmatrix} \nu_{\text{L}}^{\text{c}} \\ \nu_{\text{R}} \end{pmatrix} + \text{h.c.} \quad (5.4)$$

The relative size of m_{D} and M_{M} gives rise to different scenarios. In the following, we consider the case $M_{\text{M}} \gg m_{\text{D}}$, which is referred to as the Type I seesaw limit. This scenario is particularly interesting as it provides a mechanism to explain why the masses of the active neutrinos are exceptionally small.

In the seesaw limit, the masses of the physical neutrino states, $m_{\pm}$, defined as the eigenvalues of the mass matrix in Eq. (5.4), can be approximated as

$$m_{\pm} \approx \frac{1}{2} M_{\text{M}} \pm \frac{1}{2} \left(M_{\text{M}} + \frac{2m_{\text{D}}^2}{M_{\text{M}}} \right). \quad (5.5)$$

The masses of a light neutrino state, $\mathcal{V}$, and a heavy neutrino state, $\mathcal{N}$, can be identified as

$$|m_{\mathcal{V}}| \approx \frac{m_{\text{D}}^2}{M_{\text{M}}}, \quad (5.6)$$

and

$$m_{\mathcal{N}} \approx M_{\text{M}}. \quad (5.7)$$

As a large value of $m_{\mathcal{N}}$ results in a small value of $|m_{\mathcal{V}}|$, this process provides a mechanism that explains the smallness of the mass of the active neutrinos.

The physical states, corresponding to the eigenvectors of the mass matrix in Eq. (5.4), can be written, in the seesaw limit, as

$$\mathcal{V} \approx (\nu_{\text{L}} + \nu_{\text{L}}^{\text{c}}) - \frac{m_{\text{D}}}{M_{\text{M}}} (\nu_{\text{R}} + \nu_{\text{R}}^{\text{c}}), \quad (5.8)$$

and

$$\mathcal{N} \approx (\nu_{\text{R}} + \nu_{\text{R}}^{\text{c}}) + \frac{m_{\text{D}}}{M_{\text{M}}} (\nu_{\text{L}} + \nu_{\text{L}}^{\text{c}}). \quad (5.9)$$

Each of these fields obeys the Majorana condition, which implies that the particle state is equivalent to the antiparticle state, i.e.

$$\mathcal{V} = \mathcal{V}^{\text{c}} \quad \text{and} \quad \mathcal{N} = \mathcal{N}^{\text{c}}. \quad (5.10)$$

To make the Majorana nature of the fields explicit, we define the flavour vectors

$$v \equiv \mathcal{V} + \mathcal{V}^c, \quad (5.11)$$

and

$$N \equiv \mathcal{N} + \mathcal{N}^c. \quad (5.12)$$

On the one hand, the fields v are almost entirely left-handed, as can be deduced from Eq. (5.8) and the large value of m_D/M_M in the seesaw limit. Moreover, the couplings of v to the SM gauge bosons are essentially the same as those of the SM neutrinos. These fields are therefore identified as the active neutrinos, and take part in the unsuppressed weak interactions. Furthermore, from Eq. (5.6), we see that the seesaw mechanism implies that their mass is small. On the other hand, the fields N correspond to the heavy neutrinos, the hypothetical particles searched for in the thesis work. The properties of these particles are the following:

- They have a large mass, M_M , as inferred from Eq. (5.7).
- They are dominantly right-handed, as inferred from Eq. (5.9).
- They are almost completely sterile, as the right-handed neutrinos, ν_R , are singlets under the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetry.
- Although the N states participate in the weak interactions, the amplitudes of the corresponding process are suppressed by the active-sterile mixing matrix, θ , defined as

$$\theta \equiv \frac{m_D}{M_M}, \quad (5.13)$$

which is small in the seesaw limit.

5.3 The neutrino minimal standard model

The extension of the minimal standard model, described in Section 1.1, with three right-handed neutrinos is known as the neutrino minimal standard model (ν MSM) [93], which predicts the existence of three heavy neutrinos, noted as N_1 , N_2 , and N_3 . While the primary motivation for elaborating this model was the generation of the active neutrinos masses via the seesaw mechanism, it was found to have promising cosmological implications beyond that aim.

The Lagrangian density of the ν MSM can be formulated, with the flavour indices suppressed, as

$$\mathcal{L}_{\nu\text{MSM}} = \mathcal{L}_{\text{MSM}} + i\bar{\nu}_R \partial_\mu \gamma^\mu \nu_R - \bar{L} F \nu_R \Phi - \frac{1}{2} M_M \bar{\nu}_R^c \nu_R + \text{h.c.}, \quad (5.14)$$

where $\mathcal{L}_{\text{MSM}}$ is the Lagrangian density of the minimal SM, L are the left-handed lepton doublets, F is the matrix of Yukawa couplings, and Φ is the Higgs doublet. This model is considered as the simplest renormalisable extension of the SM in the neutrino sector, as it leaves the gauge symmetry of the SM intact and preserves its bosonic content. Moreover, the fact that three heavy neutrinos are added to the theory conserves the number of lepton flavour families, which




Figure 5.1: Illustration of the particle content of the ν MSM. The heavy neutrinos, N_1 , N_2 , and N_3 , are placed in arbitrary lepton families.

is considered as a natural choice. An illustration of the particle content in the ν MSM is shown in Fig. 5.1.

As anticipated in Section 1.5, the theoretical framework that supports the existence of heavy neutrinos can address several of the questions left open in the SM. Other than being involved in the seesaw mechanism, which generates the mass of the active neutrinos, they have two additional cosmological implications:

1. The first heavy neutrino, N_1 , with a mass in the keV range, is predicted to be a particle candidate for dark matter, provided that its lifetime is greater than the age of the universe [93].
2. The oscillations between the other two quasi mass-degenerate heavy neutrinos, N_2 and N_3 , provide a mechanism for baryogenesis, and can therefore explain the asymmetry in the abundance of matter and antimatter in the universe [94]. These neutrinos, with masses predicted to be in the GeV scale, are the particles searched for in the thesis work.

5.4 The phenomenological frame

Going a step closer to the search for heavy neutrinos presented in the next chapter, we now discuss the phenomenological considerations involved in the design of this analysis and the interpretation of its results. In particular, the framework used for the search is determined by the phenomenology described in Ref. [95], which relies on a few assumptions:

1. Except for the addition of the heavy neutrinos to the theory, the physics described by the SM remains unchanged.
2. The heavy neutrinos only interact with the SM through the mixing with the SM neutrinos, as predicted in the ν MSM. The amplitudes to the three flavour families are denoted with

V_{eN} , $V_{\mu N}$, and $V_{\tau N}$, and are defined such that

$$|V_{\ell N}|^2 \equiv \sum_i \theta_{\ell i} \theta_{\ell i}^*, \quad \ell = e, \mu, \tau, \quad (5.15)$$

where i runs over the number of HNLs and the mixing matrix θ is defined in Eq. (5.13).

3. In case there is more than one HNL in nature, they are assumed not to be mass degenerate, such that they can be kinematically distinguished by experiments. As a result, the different HNLs do not interfere, implying that the production and decay processes for a given HNL are treated independently.

A consequence of the third assumption is that the phenomenology can be fully described relative to one heavy neutrino, N , i.e. with four parameters. The parameters are typically chosen to be the mass, m_N , and the mixing angles, $V_{\ell N}$, to the three active neutrinos. Alternative parametrisations can make use of (i) the quantity $|V_N|^2$, defined as

$$|V_N|^2 \equiv |V_{eN}|^2 + |V_{\mu N}|^2 + |V_{\tau N}|^2, \quad (5.16)$$

of (ii) the mixing ratios, r_ℓ , expressed as

$$r_\ell \equiv |V_{\ell N}|^2 / |V_N|^2, \quad (5.17)$$

and of (iii) the proper mean lifetime, τ_N , which depends on m_N and $|V_N|^2$ according to

$$\tau_N \sim |V_N|^{-2} m_N^{-5}. \quad (5.18)$$

By construction, the sum of the mixing ratios to the three flavour families is constrained to one, i.e. $r_e + r_\mu + r_\tau = 1$.

The results from most of the experimental searches so far were used to constrain the scenarios in which the HNL mixes exclusively with one flavour of the active neutrino, i.e. with r_e , r_μ , or r_τ equal to 1. Although these scenarios are useful for the purposes of benchmarking, they are not well motivated from the phenomenology perspective, as they do not accommodate the oscillation phenomenon of the active neutrinos. The parametrisation in r_ℓ is thus useful, as it provides a straightforward expression for mixed-flavour mixing scenarios. As will be discussed in the next chapter, the framework used in the search for HNLs carried out in the thesis work is well suited for interpreting the results for such non-minimal scenarios.

Besides their mass and mixing pattern with the active sector, the heavy neutrinos can also be of different nature. In the context of this work, they are considered to be either Majorana or Dirac-like particles. For the Majorana case, the fact that the particle and antiparticle states are identical implies that lepton-number-violating processes are allowed. Alternatively, a Dirac-like state is obtained when two quasi-degenerate Majorana neutrinos interfere such that the lepton-number-violating decays are canceled. In this case, only lepton-number-conserving processes are allowed. This property is taken into account when interpreting the results of the search, as discussed in the next chapter.

The HNLs that are probed in this work have masses in the GeV scale. The signal particles must therefore originate from the decays of heavy particles. As far as the CMS experiment is concerned, previous searches considered the decays of W bosons as the source of HNLs. For the search designed during the thesis work, instead, a new production channel is investigated in the decays of B mesons, which have masses in the range 5.3 – 6.3 GeV. The associated decay rates are discussed in Section 6.2.1. There are several motivations for considering beauty mesons as the source of HNLs compared to W bosons:

- The B mesons are produced about a factor one thousand more abundantly in proton-proton collisions than W bosons [96]. As a consequence, they are a more prominent source of neutrinos, and potentially, of heavy neutrinos.
- The B mesons are lighter than the W bosons and their decays are therefore less boosted. Because the displacement of the decay vertex, L_{xyz} , is related to $c\tau_N$ as $L_{xyz} = \beta\gamma c\tau_N$, where $\beta\gamma$ is the boost factor, we deduce that, at equal lifetime τ_N , the particles originating from the decay of a B meson are less displaced than those originating from a W boson. As a consequence, HNLs originating from the decays of beauty mesons are more likely to be comprised within the CMS acceptance.

Before going to the actual description of the search, a few words are spent in the next section on the existing constraints on the HNL parameters obtained from other searches.

5.5 Experimental landscape

An overview of the sensitivity to heavy neutrinos from past and future experiments is presented. Thorough discussions can be found for instance in Refs. [91, 97].

First of all, there exist constraints from the Big Bang nucleosynthesis [98] and the seesaw mechanism [99]. The former exclude regions of the parameter space in which the HNLs are highly long lived, since such particles would have otherwise had an impact on the production rate of the different nuclei. The latter exclude the possibility that HNLs have such weak mixing with the active neutrinos that they would not support the observation of the neutrino oscillations.

These constraints are indicated in Fig. 5.2, which shows the sensitivity in the $|V_N|^2 - m_N$ parameter space for the case in which the Majorana HNL mixes exclusively with the muon sector, i.e. with $r_\mu = 1$. The filled areas show the exclusion limits obtained from past searches, while the dotted lines project the sensitivity of future projects or approved experiments.

Previous searches were conducted in both beam dump and collider experiments. The beam dump experiments, such as CHARM [100] or PS191 [101], have sensitivity towards lower masses compared to collider experiments, as heavy neutrinos are probed in the decays of lighter particles, such as the K or D mesons. Instead, because collider experiments can reach higher energies, the decay of heavier particles can be used as the source of neutrinos. For instance, the Belle [102] and LHCb [103] experiments consider B mesons as the signal production mode. The results of these experiments are therefore particularly interesting as they were obtained using similar




Figure 5.2: Exclusion limits on $|V_N|^2$ as a function of m_N , for the case in which the HNL mixes exclusively with the muon sector. The filled areas correspond to existing limits, while the dashed lines correspond to projections from future experiments. The figure is taken from Ref. [33].

processes as in the thesis work. On the other hand, the limits from the ATLAS [104, 105] and CMS [106, 107] experiments were obtained based on the decays of W bosons. Moreover, searches targeting scenarios in which the heavy neutrino is either prompt or long lived were conducted. While the former scenario has sensitivity towards larger masses, the latter has sensitivity towards smaller values of $|V_N|^2$, as can be inferred from Eq. (5.18).

Finally, future experiments, either approved or in the approval process, performed feasibility studies to estimate their sensitivity to HNLs. For instance, the SHiP [108] and FASER [109] experiments will target long-lived signals in the decays of heavy hadrons and are designed such that their sensitivity extends towards low values of $|V_N|^2$.

Chapter 6

Search for long-lived heavy neutrinos in the decays of B mesons

A search for long-lived, GeV-scale heavy neutrinos produced in the leptonic or semileptonic decays of B mesons is performed using proton-proton collision data collected at a centre-of-mass energy of 13 TeV. The search exploits a beauty-hadron-enriched event sample corresponding to an integrated luminosity of 41.6 fb^{-1} . In the following, heavy neutrinos are denoted with the symbol N and are also referred to as heavy neutral leptons (HNLs).

Searching for heavy neutrinos in beauty meson decays is a novel approach at CMS. Indeed, these hypothetical particles were so far probed solely in the decays of W bosons [106, 107, 110–113]. As discussed in Section 5.4, decays of B mesons are a promising production channel, notably because these hadrons are produced significantly more abundantly in proton-proton collisions than W bosons, and thus provide a stronger source of neutrinos, and potentially of heavy neutrinos. For a search in this novel production channel to be significant, a large number of decays from B mesons must be analysed. This ability is provided by the collection, in 2018, of the B-parking data set, presented in Section 2.3.1, which contains of order 10^{10} $b\bar{b}$ events with a high purity.

The analysis takes as a framework the phenomenology described in Ref. [95] and summarised in Section 5.4. In particular, the HNL signal may either be a Majorana or Dirac-like particle, and is parametrised with four parameters, typically chosen to be the mass, m_N , as well as the mixing amplitudes, $V_{\ell N}$, to the three lepton flavour families, $\ell = (e, \mu, \tau)$. Other possible parametrisations make use of the quantity $|V_N|^2$, defined in Eq. (5.16), the mixing ratios, r_ℓ , defined in Eq. (5.17), and the lifetime, τ_N , which depends on m_N and $|V_N|^2$ according to the relation in Eq. (5.18). The lifetime, τ_N , and its conversion in units of length, $c\tau_N$, are used interchangeably throughout the chapter.

The present search probes the following signal processes, in either the semileptonic or leptonic B decay channels:

$$\begin{aligned} B &\rightarrow \ell_B N X; & N &\rightarrow \ell^\pm \pi^\mp && \text{(Semileptonic B decays),} \\ B &\rightarrow \ell_B N; & N &\rightarrow \ell^\pm \pi^\mp && \text{(Leptonic B decays),} \end{aligned} \tag{6.1}$$




Figure 6.1: Feynman diagrams representing the semileptonic (*upper*) and leptonic (*lower*) B meson decay into a charged lepton, ℓ_B , a hadronic system, X , in case of the semileptonic decay, and a neutrino, ν_{ℓ_B} , which contains a small admixture of a heavy neutrino, N . The N mass eigenstate propagates, and according to its admixture of the neutrino flavour eigenstate, ν_ℓ , decays into a charged lepton, $\ell^\pm$, and charged pion, $\pi^\mp$.

where

- The B meson may either be $B^\pm$, B^0 , B_s^0 , or $B_c^\pm$, denoted hereafter as B_u , B_d , B_s , and B_c , respectively. The charge conjugate states are also considered.
- For the semileptonic B decays, the hadronic recoil system, X , is a meson that is not reconstructed. It follows that the B meson decays are treated inclusively.
- The charged lepton, ℓ_B , arising from the B decay is either a muon or an electron. It is referred to, in the context of this work, as the *primary* lepton.
- N is the hypothetical long-lived heavy neutrino. It decays exclusively into a displaced charged lepton, $\ell^\pm$, which is either a muon or an electron, and a displaced charged pion, $\pi^\mp$. The decay is required to occur within the silicon tracker volume, defined as the region with $L_{xy} < 1$ m, where L_{xy} is the displacement in the transverse plane between the HNL decay vertex and the beamspot.

The leading order Feynman diagrams of the corresponding processes are shown in Fig. 6.1.

The search probes HNLs with masses in the range $1 < m_N < 6$ GeV. The upper bound is imposed by the mass of the B_c meson, which is the heaviest meson with mass of about 6.3 GeV,

while the other beauty mesons have a mass about 1 GeV smaller. Moreover, heavy neutrinos with decay lengths approximately in the range $10^{-2} < c\tau_N < 10^4$ mm are considered.

The assumption that the HNL is a long-lived particle is a key feature of the analysis: based on Eq. (5.18), we see that large lifetimes give access to low values of $|V_N|^2$, which is a region of the parameter space that remains largely unexplored experimentally. However, probing signatures of a long-lived particle decay is a challenging task at CMS [114], as the detector was primarily designed to capture signatures from promptly decaying particles. Multiple strategies are therefore put in place to maximise the sensitivity of the search. Among them, a signal-to-background multivariate discriminator based on parametric neural networks [115] is implemented to optimise the performance of the search over a wide range of signal hypotheses. Moreover, the events are distributed into different search categories to enhance the sensitivity of the analysis to different signal signatures.

The final states analysed in this search contain three low- p_T objects: (i) the charged lepton from the B meson decay, ℓ_B , (ii) the charged lepton from the N decay, $\ell^\pm$, and (iii) the charged pion from the N decay, $\pi^\mp$. The displaced lepton and pion are required to originate from a common vertex, referred to as the displaced vertex (DV), away from the beamspot with a sizable distance. Furthermore, the two leptons, ℓ_B and $\ell^\pm$, can either be a muon or an electron, with the condition that one of them is a muon that triggers a B-parking line. This condition gives rise to three possible final states, based on the respective flavour of the leptons: $(\ell_B, \ell^\pm) = (\mu\mu, \mu e, e\mu)$. The $\mu\mu$ channel is referred to as the dimuon channel, while the μe and $e\mu$ possibilities constitute the mixed-flavour channel.

As the heavy neutrino is assumed to decay exclusively into visible objects and within the tracker volume, its invariant mass can be fully reconstructed. The strategy of the search thus is that of a *bump hunt*, implying that a peak compatible with the expected signal shape is searched for in the invariant mass distribution of the displaced $\ell^\pm\pi^\mp$ system. The SM backgrounds, primarily arising from soft QCD processes, passing the event selection are estimated from a direct fit to the data. The signal is finally extracted from the data using a maximum likelihood fit to the $\ell^\pm\pi^\mp$ mass distributions performed simultaneously in all the search categories.

This search explores over one hundred signal mass hypotheses with an unprecedented resolution for an HNL search, stemming from the bump hunt strategy. Consequently, the mass of the new particle could be measured, in case a signal were to be found in data. Moreover, multiple signal scenarios are considered. While most of the searches so far [97] considered scenarios in which the HNL mixes exclusively with one lepton family, i.e. with two out of three mixing ratios, r_ℓ , set to zero, this search explores as well mixed-flavour mixing scenarios, with at least two non-zero mixing ratios. Constraints on such scenarios are of great interest to the phenomenology community [116, 117]. Moreover, the results are interpreted for both the Majorana and Dirac-like cases.

All aspects of the analysis were initiated, designed, and validated in the context of the thesis work. Results were primarily derived for the dimuon channel, while members of Italian institutes at INFN Rome and Sapienza derived the results for the mixed-flavour channel, using the same methods established for the analysis in the dimuon channel. The final combination of the two

channels to extract the final results, as well as the coordination of the effort towards publication were also an integral part of the thesis work.

The results from this research project were presented at the Moriond/QCD 2024 conference [118] and were published in the *Journal of High Energy Physics* the same year [119]. Moreover, this work is summarised in two review papers to be published in the *Physics Reports* journal, the first one summarising the CMS data parking techniques [34], and the second one giving an overview of the CMS HNL searches [120].

The chapter is organised as follows. A description of the B-parking data set is presented in Section 6.1. Section 6.2 gives an overview of the main features related to the simulated signal, namely the HNL production and decay widths, the signal sample generation scheme, and the expected signatures from a signal event. The event reconstruction is discussed in Section 6.3 and the event selection and categories are described in Section 6.4. Corrections applied on the simulated events are summarised in Section 6.5. The signal yield normalisation and shape parametrisation are presented in Section 6.6, and the estimation of the background is discussed in Section 6.7. Sources of systematic uncertainties are described in Section 6.8 and the results are presented in Section 6.9. Finally, the main aspects of the search are summarised in Section 6.10.

6.1 The B-parking data sample

Collected in 2018 using the data parking strategy described in Section 2.3.1, the B-parking data set is the largest sample of $b\bar{b}$ events ever collected by the CMS experiment [34, 35]. The dedicated triggers, described in the following, accumulated approximately 1.2×10^{10} events, out of which about 1.0×10^{10} events originate from the decay of a beauty hadron. The corresponding integrated luminosity is 41.6 fb^{-1} . The design of the trigger strategy took place before the start of the PhD and was therefore not part of the thesis work.

Both the L1 trigger and HLT algorithms were designed to capture events arising from the leptonic and semileptonic decays of B hadrons. The typical signature of these decays is the presence of a lepton which (i) has low transverse momentum and (ii) whose track is displaced with respect to the luminous region, since B hadrons have a relatively long lifetime. Because of the excellent ability to reconstruct and identify muons at CMS, these particles are the ones used to acquire the B-parking data sample.

A series of L1 trigger algorithms are designed such that they select events with a muon, provided that it has a minimum p_T ranging from 7 to 12 GeV, depending on the algorithm, and that it is centrally produced, with $|\eta| < 1.5$. At the HLT, the muons are not only reconstructed using the information from the muon spectrometer, as for L1 muons, but also using the information from the tracking system, as explained in Section 2.3. Therefore, the transverse impact parameter significance, $|d_{xy}|/\sigma_{d_{xy}}$, of their track can be reconstructed. This quantity, defined as the ratio of the transverse impact parameter, $|d_{xy}|$, to the associated uncertainty, $\sigma_{d_{xy}}$, is useful to quantify the displacement of a track with respect to the luminous region. A displaced track may arise from the decay of a B hadron, stemming from its sizable lifetime. For that reason, the HLT paths additionally impose a minimum threshold on $|d_{xy}|/\sigma_{d_{xy}}$ to maximise the probability

Table 6.1: List of the HLT trigger conditions. Each line represents the combination of the minimal thresholds applied on the muon p_T and $|d_{xy}|/\sigma_{d_{xy}}$ for a given HLT path.

HLT μ p_T (GeV)	HLT μ $ d_{xy} /\sigma_{d_{xy}}$
7	4
8	3
8	5
8	6
8.5	3.5
9	4
9	5
9	6
10.5	3.5
12	6

to select, within the constraint of the maximum allocated trigger rate, a displaced muon arising from the decay of a B hadron. The list of the HLT trigger paths used to collect the B-parking data set is given in Table 6.1.

A feature of the collection of the B-parking data sample is that the trigger thresholds were not kept constant over time and were, instead, progressively relaxed throughout an LHC fill. Indeed, the instantaneous luminosity, L , decreases during a fill, as a consequence of the losses from beam interactions. Therefore, the rate of a given HLT path decreases with time as well. To exploit the trigger system at its full capacity, the thresholds for both the L1 trigger and the HLT were progressively lowered such that the consequent increase in the rate compensates the loss induced by the decreasing luminosity. With this strategy, a large amount of data could be collected, without perturbing the data taking of the regular stream. The evolution of the trigger rate as a function of time for both the L1 trigger and HLT is shown in Fig. 6.2 for a representative LHC fill. We verify that the change in the trigger thresholds, indicated with the vertical lines, leads to an increased trigger rate, such that the trigger system operates close to its design limit throughout the entire fill. The fact that the triggers were not always active gives rise to a complex trigger structure. A challenge in the analysis is to reproduce the associated effect for simulated events, as discussed in Section 6.5.

Finally, the conservation of the beauty quantum number in strong processes imposes that the B hadrons are produced in pairs in an event. The hadron decaying into the muon responsible for the positive trigger decision is referred to as the *tag* hadron, while the other one, left unbiased, is referred to as the *probe* hadron. In the context of this analysis, the tag B meson is the one that is considered to decay into a heavy neutrino.

6.2 Simulated signal samples

Signal events are simulated to optimise the event selection criteria, to extract the signal shape of the $\ell^\pm\pi^\mp$ invariant mass distributions used to fit the data, and to obtain the signal efficiencies needed to normalise the signal events. As no pre-existing tool allowed for the simulation of heavy




Figure 6.2: B-parking rates of the L1 trigger (*left*) and HLT (*right*), and number of pileup vertices as a function of time for an arbitrary LHC fill. The modification of the trigger thresholds are indicated with the vertical lines. For the HLT, the black markers show the rate for the promptly reconstructed stream, while the blue ones show the rate for the B-parking stream. The figures are taken from Ref. [34].

neutrinos in beauty mesons decays, a crucial part of the work was the design of a framework making the simulation of such events possible for given values of m_N and $c\tau_N$. In the following, a discussion on the HNL production and decay widths is presented in Section 6.2.1, the signal event sample generation is described in Section 6.2.2, and the main signatures of the signal events are summarised in Section 6.2.3. We note that simulated background event samples are not needed in the analysis, because the event selection is optimised using data, as described in Section 6.4, and because the background shape is fitted directly to data, as discussed in Section 6.7.

6.2.1 Heavy neutrino production and decay widths

The capability to compute the heavy neutrino production and decay rates for the different signal hypotheses, parametrised in the following with m_N , $|V_N|^2$, and $\vec{r} = (r_e, r_\mu, r_\tau)$, is an important aspect of the analysis for two reasons: (i) the rates are used to configure the simulation of the signal samples, and (ii) they are necessary to compute the expected number of signal events for the different signal hypotheses, as described in Section 6.6.1. The HNL production and decay rates are computed following the analytical expressions presented in Ref. [95]. Based on these expressions, a framework was created in the context of the thesis to compute the production and decay rates as functions of m_N , $|V_N|^2$, and $\vec{r}$.

Heavy neutrino production

With respect to the production, the leptonic and semileptonic decays of B mesons are considered as the source of heavy neutrinos in the processes targeted in the search. Decays from B baryons are not included, as their contribution to the sensitivity of the analysis is expected to be significantly smaller than that from B mesons, resulting from the smaller fragmentation fractions to B baryons than to B mesons [3].

Table 6.2: List of the main heavy neutrino production decays from the B_u , B_d , B_s , and B_c mesons. The decay modes are categorised into leptonic processes, semileptonic processes involving a pseudoscalar meson, and semileptonic processes involving a vector meson. The dash symbol “–” indicates that there is no corresponding decay mode.

Process	B_u decays	B_d decays	B_s decays	B_c decays
Leptonic	$B_u \rightarrow \ell N$	–	–	$B_c \rightarrow \ell N$
Semileptonic (pseudoscalar)	$B_u \rightarrow \ell N D^0$	$B_d \rightarrow \ell N D$	$B_s \rightarrow \ell N D_s$	–
	$B_u \rightarrow \ell N \pi^0$	$B_d \rightarrow \ell N \pi$	$B_s \rightarrow \ell N K$	–
Semileptonic (vector)	$B_u \rightarrow \ell N D^{0*}$	$B_d \rightarrow \ell N D^*$	$B_s \rightarrow \ell N D_s^*$	–
	$B_u \rightarrow \ell N \rho^0$	$B_d \rightarrow \ell N \rho$	$B_s \rightarrow \ell N K^*$	–

The strategy employed in the analysis consists in treating the decays of the B mesons into an HNL in an inclusive manner, i.e. by considering all possible hadronic systems, labelled with X in Eq. (6.1), instead of restricting the search to a particular one. Although the invariant mass of the B mesons cannot be reconstructed for each of the HNL mass hypotheses with this strategy, the inclusive approach significantly increases the signal coverage, and hence the sensitivity of the search.

The list of the most important decays of the B_u , B_d , B_s , and B_c mesons into an HNL is given in Table 6.2. The possible decay modes are categorised into three classes of processes: (i) two-body leptonic decays, (ii) three-body semileptonic decays into a pseudoscalar meson, and (iii) three-body semileptonic decays into a vector meson. Decays involving more than one meson in the final state only contribute marginally to the total decay rate and are thus not included.

Figure 6.3 shows examples of branching fractions of three B_u decays. The branching fractions are shown for $|V_N|^2 = 1$ and for the scenario in which the HNL mixes exclusively with the electron neutrino, i.e. $(r_e, r_\mu, r_\tau) = (1, 0, 0)$, to allow for a direct comparison with the results presented in Ref. [95]. The excellent agreement between the results extracted from this reference and the ones reproduced in this work validates the framework created to compute the HNL production rates.

Additionally, we see in Fig. 6.3 that the semileptonic rate for $B \rightarrow \ell N X$ drops to zero for HNL masses $m_N > |m_B - m_X|$, where m_B and m_X are the mass of the B meson and the accompanying meson, X, respectively. Instead, HNL masses can be probed up to the mass of the mother B meson when the latter produces the heavy neutrino via the leptonic channel. This behaviour is confirmed in Fig. 6.4 (left), which shows the rate of leptonic and semileptonic decays as a function of m_N for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ and $|V_N|^2 = 1$. We see that heavy neutrinos with masses $m_N < 2$ GeV originate predominantly from semileptonic decays, while heavy neutrinos with masses $m_N > 4$ GeV originate predominantly from leptonic decays. The transition between the dominance of the semileptonic or leptonic channels takes place for HNLs with mass around 3 GeV, which will have an impact on the final results, as discussed in Section 6.9.

Finally, Fig. 6.4 (right) shows, as a function of m_N , the total decay branching fractions for




Figure 6.3: Branching fractions of the leptonic process $B_u \rightarrow eN$, the semileptonic process into a pseudoscalar meson $B_u \rightarrow eND^0$, and the semileptonic process into a vector meson $B_u \rightarrow eN\rho^0$ shown as a function of the mass m_N . The branching fractions correspond to the scenario $(r_e, r_\mu, r_\tau) = (1, 0, 0)$ and $|V_N|^2 = 1$. The coloured curves are obtained with the tool implemented in this work and the black markers are points extracted from the results presented in Ref. [95].

Table 6.3: Values of the fragmentation fractions f_q ($q = u, d, s, c$) for the different B meson species and their corresponding reference. The fragmentation fraction f_q is defined as the probability for a b quark to hadronise into a B_q meson.

Species	Fragmentation fraction	Reference
B_u	$f_u = 4.00 \times 10^{-1}$	[3]
B_d	$f_d = 4.00 \times 10^{-1}$	[3]
B_s	$f_s = 1.00 \times 10^{-1}$	[3]
B_c	$f_c = 2.61 \times 10^{-3}$	[121]

the four B mesons species, multiplied by the respective fragmentation fractions, f_q , listed in Table 6.3, for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ and $|V_N|^2 = 1$. For HNL masses $m_N \gtrsim 3$ GeV the dominant contribution comes from B_c decays, despite the much smaller fragmentation fraction of this meson compared to the others. This behaviour is exploited in the analysis, notably with the design of a dedicated event categorisation based on $m(\ell_B \ell^\pm \pi^\mp)$, as discussed in Section 6.4.2.

Heavy neutrino decay

The decays of the heavy neutrino are mediated either by charged-current or neutral-current interactions, as listed in Table 6.4. In particular, the HNL may decay either through the hadronic, leptonic, or invisible channels. The branching fractions for the three channels are shown in Fig. 6.5 as a function of m_N for the mixing scenario $(r_e, r_\mu, r_\tau) = (1/3, 1/3, 1/3)$. The hadronic channel is the dominant decay mode of the HNL. The fractions are computed based on the work




Figure 6.4: (Left) Branching fraction of the leptonic and semileptonic decays divided by the total branching fraction, shown as a function of m_N . The fractions are computed for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ and $|V_N|^2 = 1$. (Right) Branching fractions as a function of m_N for $B_q \rightarrow \mu NX$ decays, with $q = (u, d, s, c)$, multiplied by the corresponding fragmentation fraction, f_q , listed in Table 6.3. The results are shown for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ and $|V_N|^2 = 1$.

Table 6.4: List of charged-current-mediated and neutral-current-mediated heavy neutrino decays in the hadronic, leptonic, and invisible channels. The quarks u and d can have any of the up-type or down-type quark flavour, respectively, and the quark q can have any flavour. The indices α and β run over the three lepton flavours (e, μ, τ). The dash symbol “–” indicates that there is no corresponding decay mode. In the signal process, the charged-current-mediated hadronic decay into $\ell\pi$ is targeted.

	Charged-current	Neutral-current
Hadronic	$N \rightarrow \ell_\alpha^+ u \bar{d}$	$N \rightarrow \nu_\alpha q \bar{q}$
Leptonic	$N \rightarrow \ell_\alpha^+ \ell_\beta^- \bar{\nu}$	$N \rightarrow \nu_\alpha \ell_\beta^+ \ell_\beta^-$
Invisible	–	$N \rightarrow \nu_\alpha \nu_\beta \bar{\nu}_\beta$

presented in Ref. [95], and can be well reproduced with the framework designed to compute the decay rates, as can be seen in this figure.

Figure 6.6 shows, as a function of m_N , the rates, Γ_α , of HNL decays that involve a charged or neutral lepton of flavour $\alpha = (e, \mu, \tau)$, computed for the case $|V_N|^2 = 1$. The corresponding rates are referred to as the *reduced decay rates* and are denoted as $\tilde{\Gamma}_\alpha$, to be consistent with the notation used in Section 6.6. The rates are shown for the Majorana scenario, in which the charged conjugated processes must be taken into account. The decay rates for the Dirac-like scenario are therefore a factor two smaller than those in the Majorana scenario. We note that the decay rate to the tau sector is significantly smaller than the decay rates to the electron and muon sectors, owing to the larger mass of the tau lepton with respect to the mass of the electron and the muon.

The search probes heavy neutrinos decaying into $\ell\pi$, as this decay channel is the one with




Figure 6.5: Branching fraction as a function of m_N for a heavy neutrino decaying via the hadronic, leptonic, and invisible channels. The contributions from the charged-current-mediated and neutral-current-mediated decays are summed. The widths are computed for the mixing scenario $(r_e, r_\mu, r_\tau) = (1/3, 1/3, 1/3)$. The coloured curves are obtained with the tool implemented in this work and the black markers are points extracted from the results presented in Ref. [95].

the largest branching fraction among those involving uniquely visible objects. The rate of the $N \rightarrow \ell\pi$ decay is shown in Fig. 6.6 for the case in which the HNL mixes exclusively with muon neutrinos. The corresponding branching fraction ranges from about 22% to 0.5% for $m_N = 1.0$ and 6.0 GeV, respectively, as seen in the lower panel of the figure.

6.2.2 Event generation

A signal generation scheme was designed to allow for the simulation of the signal for the different m_N and $c\tau_N$ hypotheses. Two sets of signal samples are generated: the first set corresponds to the production of HNLs in the decays of B_u , B_d , and B_s mesons, while the second one corresponds to the production of HNLs in the decays of B_c mesons. The simulation of B_c decays requires a dedicated generator that produces B_c mesons directly, because the low fragmentation fraction of this meson makes the standard generation via $b\bar{b}$ pairs not computationally viable.

For the (B_u, B_d, B_s) samples, the proton-proton collisions and hadronisation are simulated with PYTHIA 8.230 [74], using the SoftQCD:nonDiffractive process to reproduce minimum bias interactions. The B mesons are generated from $b\bar{b}$ pairs according to the fragmentation fractions listed in Table 6.3. Events containing a B meson are then passed to a version of EVTGEN 1.30 [122] adapted to allow the required B meson and HNL decays: the B meson is decayed to one of the channels listed in Table 6.2 according to the respective branching fraction at a given m_N , and the HNL is decayed to $\ell\pi$. The decays are performed using a phase space decay model (PHSP). In case more than one B meson species is generated in an event, only one is chosen by EVTGEN to be decayed into an HNL. Because this choice is random, the signal events




Figure 6.6: The upper panel shows the reduced HNL decay widths as a function of m_N . The indigo, red, and orange coloured solid curves show the widths obtained for the case in which the HNL decay involves a charged or neutral electron, muon, or tau lepton, respectively. The black dashed curve corresponds to the width of the $N \rightarrow \mu\pi$ decay, computed for the case where the HNL mixes exclusively with the muon sector. All the widths are shown for the Majorana scenario, and are hence a factor two larger than the widths for the Dirac-like scenario. The lower panel shows as a function of m_N the branching fraction of the $N \rightarrow \mu\pi$ decay, for the case in which the HNL mixes exclusively with the muon sector.

must be normalised separately for each B meson species, as described in Section 6.6.1, such as to not distort the relative branching fractions presented in Section 6.2.1.

For the B_c samples, the proton-proton collisions are first generated with the BCVEGPy 2.2b generator [123], with one B_c meson per event. The events are then passed to PYTHIA and EVTGEN for the hadronisation and signal decay, in a fashion similar to that used for the (B_u , B_d , B_s) samples.

For each set of signal samples, the generated events are then propagated through the CMS detector using the GEANT4 package [75]. The presence of pileup events is reproduced in the samples with a distribution similar to the one in the B-parking data set via the simulation of minimum bias events. Finally, events are reconstructed using the same CMS software as for collision data.

A series of filters is applied to select events only within the detector acceptance, and hence reduce the size of the simulated signal samples. These filters are

- Filter on the presence of a B meson in the event.
- Filter on the presence of two leptons in the event. One lepton must be a muon with $p_T > 6.8$ GeV and $|\eta| < 1.55$ to meet the trigger requirements and the other lepton can

either be a muon or an electron, provided that it has $p_T > 1$ GeV and $|\eta| < 2.45$. The triggering muon is required to originate from a B meson or an HNL, and the invariant mass of the two leptons is constrained to be less than 10 GeV to reject events in which the HNL originates from the probe side.

- Filter on the presence of a charged pion originating from the decay of the HNL and with $p_T > 0.5$ GeV.
- Filter requiring the transverse displacement, L_{xy} , to be smaller than 1.3 m, to enforce the HNL decay to take place within the tracker volume.

The efficiency of these filters ranges from approximately 10^{-3} to 10^{-2} , depending on the m_N and $c\tau_N$ hypotheses. Furthermore, the efficiency is obtained separately for the different B meson species, and for the dimuon and mixed-flavour channels. Finally, the rate of events on the tag side is approximately 90%.

The signal samples are produced for 103 mass hypotheses in the range $1 < m_N < 6$ GeV, chosen with a step size corresponding to approximately twice the detector resolution, σ , discussed in Section 6.6.6. The (B_u , B_d , B_s) samples are generated for masses $1 < m_N < 5$ GeV, and the B_c samples for masses $3 < m_N < 6$ GeV, according to Fig. 6.4 (right). For each value of m_N , one to four samples are generated with different $c\tau_N$ hypotheses in the range $0.1 < c\tau_N < 1000$ mm. For the samples in the dimuon channel, the events are generated for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$, while for the samples in the mixed-flavour channel, they are generated for the mixing scenario $(r_e, r_\mu, r_\tau) = (1/2, 1/2, 0)$. In both cases, the HNL is assumed to be a Majorana particle. As described in Section 6.6.2, 6.6.3, and 6.6.4, weights are applied on the events of these samples to interpret the results for different mixing scenarios, for the Dirac-like case, and for different $c\tau_N$ hypotheses, respectively.

6.2.3 Signal signatures

Examples of the distributions of key variables are shown in Fig. 6.7 using MC events at generator-level, i.e. without the simulation of the CMS detector and without the reconstruction of the particles with the CMS software. The distributions of $p_T(\ell_B)$, $p_T(\ell^\pm)$, $m(\ell_B \ell^\pm \pi^\mp)$, and L_{xy} are presented for the signal hypotheses (i) $m_N = 1$ GeV, $c\tau_N = 1000$ mm, (ii) $m_N = 3$ GeV, $c\tau_N = 100$ mm, and (iii) $m_N = 4.5$ GeV, $c\tau_N = 1$ mm. These hypotheses are close to the expected sensitivity of the search for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$. Events are taken from the (B_u , B_d , B_s) simulated samples, and the corrections discussed in Section 6.5 are not applied.

The three signal hypotheses yield significantly different shapes of the distribution for each quantity. This behaviour shows that a unique set of event selection criteria would not be optimal across the wide range of m_N and $c\tau_N$ values investigated in the search. The solution designed to obtain a signal-to-background discriminator with good performance for all signal hypotheses is discussed in Section 6.4.3.

The distributions of the transverse momentum of the primary and displaced leptons are




Figure 6.7: Distribution of the primary lepton p_T (upper left), displaced lepton p_T (upper right), invariant mass of the $\ell_B^\pm \pi^\mp$ system (lower left), and L_{xy} of the DV (lower right). The distributions are shown for generator-level events from the (B_u , B_d , B_s) samples simulated with (i) $m_N = 1$ GeV, $c\tau_N = 1000$ mm, (ii) $m_N = 3$ GeV, $c\tau_N = 100$ mm, and (iii) $m_N = 4.5$ GeV, $c\tau_N = 1$ mm. No corrections are applied on the simulated events. The distributions are normalised to unit area.

shown in the upper left and right plots of Fig. 6.7, respectively. These distributions show a step behaviour, which is explained by the presence of the triggering muon in the event. As summarised in Table 6.1, the lowest threshold on the muon p_T among the B-parking triggers equals 7 GeV. Therefore, depending on whether it is the primary or the displaced muon that triggers, the respective p_T distributions will preferably populate the bins above 7 GeV. For signal events with low values of m_N , the triggering muon is preferably the primary muon, while for signal events with large values of m_N , the triggering muon is instead the displaced muon. The addition of events in which the displaced muon triggers was found to improve the sensitivity of the search by up to a factor of about 3.

Figure 6.7 (lower left) shows the invariant mass of the $\ell_B^\pm \pi^\mp$ system. We verify that the mass of the mother B meson is not fully reconstructed for signal hypotheses with low values

of m_N , for which the semileptonic B decays dominate, according to Fig. 6.4 (left). For signal hypotheses with larger masses, the mother B meson decays predominantly via the leptonic channel, and the distribution hence presents a peak at the mass of the B meson. For such events, the quantity $m(\ell_B \ell^\pm \pi^\mp)$ offers an efficient discrimination power against background events, as discussed in Section 6.4.

Finally, Fig. 6.7 (lower right) shows the distribution of the transverse displacement of the DV, L_{xy} . As opposed to the previous quantities, this variable does not depend on the value of m_N but rather on the value of $c\tau_N$. The larger is the value of $c\tau_N$, the more displaced is the DV with respect to the luminous region. As described in Section 6.4, the displacement-related quantities are also used to discriminate between signal and background events.

6.3 Event reconstruction

The presence of heavy neutrinos in the data is probed in final states containing a primary charged lepton, ℓ_B , a displaced charged lepton, $\ell^\pm$, and a displaced charged pion, $\pi^\mp$. As the primary lepton originates from the decay of a B meson, which has a relatively long lifetime, it too may be displaced with respect to the proton-proton collision vertex. The reconstruction of the $\ell_B \ell^\pm \pi^\mp$ signal candidates is discussed in the following.

The charged leptons, ℓ_B and $\ell^\pm$, can either be muons or electrons. Tau leptons are not considered in this work because their short lifetime [58] would either suppress the sensitivity of this channel or require dedicated analysis treatments to be accounted for.

The particles are reconstructed with the standard CMS algorithms, described in Section 2.4. Further conditions are imposed on the identification of the leptons, as discussed in Section 6.4.1. An alternative muon reconstruction algorithm [124] that does not exploit tracker information and does not impose a track displacement criteria with respect to the beamspot was studied in view of improving the reconstruction efficiency of muon candidates. However, it was found not to improve the sensitivity of the search and was thus eventually not used for the analysis. Finally, the charged pions are reconstructed as tracks to which the pion mass is assigned.

The tracks of the displaced objects, $\ell^\pm$ and $\pi^\mp$, are required to originate from a common displaced vertex (DV), which is reconstructed using a kinematic vertex fitter [125]. This algorithm is based on a least mean square minimization with Lagrange multipliers and Kalman filter techniques. The momenta of the displaced tracks are recomputed using the DV as a constraint, to improve the resolution in mass of the reconstructed $\ell_B \ell^\pm \pi^\mp$ candidates.

Another characteristic of the search is the presence of a triggering muon, which can either be ℓ_B or $\ell^\pm$, in the final state. A muon is considered to be triggering if it is matched to a trigger object corresponding to a B-parking trigger line. The matching consists in requiring that (i) the muon fired the given trigger object, that (ii) the muon and trigger object are close to each other within $\Delta R < 0.05$, and that (iii) the relative difference in their transverse momentum is less than 10%.

For simulated signal events, the reconstructed objects are additionally matched to generator-




Figure 6.8: Efficiency of the $\ell_B \ell^\pm \pi^\mp$ candidate reconstruction shown in bins of L_{xy} and $p_T(\mu^\pm \pi^\mp)$ for events in the dimuon channel simulated with $m_N = 3$ GeV and $c\tau_N = 2$ m.

level objects, to prevent the unwanted reconstruction of events arising from combinatorial processes. A reconstructed object is matched to a generator-level object if the two objects are within $\Delta R < 0.25$ from each other and if the relative difference in their transverse momentum is less than 25% for the leptons and 50% for the pion. Moreover, the mother particle of the generator-level ℓ_B must be a B meson, while the mother particle of the generator-level $\ell^\pm$ and $\pi^\mp$ must be an HNL. This matching has an efficiency of 80% that is accounted for when computing the expected number of signal events, as discussed in Section 6.6.1.

In either data or MC, the signature of the final state consists in three objects with low transverse momenta (following from the HNL production in B meson decays), two of these objects being displaced (following from the requirement that the HNL is long lived). The combination of these features has a direct impact on the $\ell_B \ell^\pm \pi^\mp$ candidate reconstruction efficiency. The CMS detector was indeed primarily designed to study final states with prompt objects, i.e. originating close to the beamspot and having medium to large transverse momenta.

The reconstruction efficiency is shown in Fig. 6.8 in bins of the transverse displacement, L_{xy} , of the DV with respect to the beamspot, and in bins of the p_T of the $\mu^\pm \pi^\mp$ system. The efficiency is obtained using signal events in the dimuon channel simulated with $m_N = 3$ GeV and $c\tau_N = 2$ m, and is computed as the ratio of the number of reconstructed $\mu_B \mu^\pm \pi^\mp$ candidates to the number of generated candidates that are within the detector acceptance. The $\mu_B \mu^\pm \pi^\mp$ candidate reconstruction efficiency significantly decreases with larger values of L_{xy} , reaching only a few percent for $L_{xy} > 30$ cm. Signal hypotheses with long lifetimes will be the most impacted. Moreover, the efficiency decreases with lower values of p_T . To push the sensitivity of the search to a competitive level, dedicated techniques were designed to optimise the signal significance, as discussed in the next section.

6.4 Event selection and categories

After the $\ell_B \ell^\pm \pi^\mp$ candidates are reconstructed in data, only those likely to originate from an HNL decay are selected. The selection process takes place in three steps. In the first step, described in Section 6.4.1 and referred to as *preselection*, inclusive thresholds on the value of kinematic and topological variables describing the event are imposed in such a way that the overwhelming background is rejected and that the majority of the signal events across the large range of m_N and $c\tau_N$ values is preserved. In the second step, the events passing the preselection conditions are categorised according to the criteria defined in Section 6.4.2. Finally, in the third and last step, presented in Section 6.4.3, the events are selected if they pass a given threshold on the output of a multivariate algorithm, trained in the different categories.

6.4.1 Preselection

In a single event, there can be of order a hundred $\ell_B \ell^\pm \pi^\mp$ candidates reconstructed in data following the description presented in Section 6.3. The vast majority of these candidates originates from background processes that give rise to final states with significantly different signatures than those from signal processes. Therefore, a series of criteria identifying the main properties of the signal must be imposed to select signal-like candidates and suppress the ones coming from background. This analysis step, referred to as the preselection, is designed in such a way that the genuine signal candidates, obtained from the simulated samples described in Section 6.2, are mostly retained over the wide range of masses and lifetimes probed in the search.

The preselection consists in applying cuts, i.e. minimum or maximum thresholds, on the value of variables describing the event and offering a good discrimination between signal and background processes. For each variable, listed below, the preselection cut is optimised against the efficiency, defined as the ratio of the number of reconstructed candidates obtained with and without applying the cut. While the value of the threshold must be chosen such that the background efficiency is as low as possible, it should however not be too stringent, to preserve a high signal efficiency for the different signal hypotheses.

For the background candidates, a small portion of the B-parking data set corresponding to less than 0.01% of the events is used. For the signal candidates, the preselection cuts are optimised using simulated events for three benchmark signal scenarios:

1. $m_N = 1$ GeV, $c\tau_N = 1000$ mm.
2. $m_N = 3$ GeV, $c\tau_N = 100$ mm.
3. $m_N = 4.5$ GeV, $c\tau_N = 1$ mm.

These scenarios are chosen as they provide a good representation of the range of possible kinematics and lifetimes in a region of the parameter space close to the expected sensitivity of the search.

Typical variables providing a good discrimination between signal and background candidates and constrained in the preselection are




Figure 6.9: Distributions of the displaced pion p_T (left) and $|d_{xy}|/\sigma_{d_{xy}}$ (right) in data and in simulated events from the (B_u , B_d , B_s) samples corresponding to the three benchmark signal hypotheses. The distributions are shown for events in the dimuon channel and are normalised to unit area. The last bin shows the overflow. The corresponding preselection conditions are $p_T(\pi) > 1$ GeV and $|d_{xy}|/\sigma_{d_{xy}}(\pi) > 5$, respectively.

- Variables describing the kinematics of the final state particles. An example is given in Fig. 6.9 (left), which shows the distribution of the displaced pion p_T for the background and the three benchmark signal hypotheses, in the dimuon channel. The background, which arises mainly from soft QCD processes, as discussed in Section 6.7.1, has a softer p_T distribution compared to that of the signal.
- Variables describing the displacement of the track of the final state particles with respect to the beamspot. Such variables are the transverse and longitudinal impact parameters, $|d_{xy}|$ and $|d_z|$, as well as their significance, defined as $|d_{xy}|/\sigma_{d_{xy}}$ and $|d_z|/\sigma_{d_z}$, where $\sigma_{d_{xy}}$ and σ_{d_z} are the associated uncertainties. The transverse impact parameters are computed with respect to the beamspot, while the longitudinal ones are computed with respect to the primary vertex, defined in Section 2.4.1. Figure 6.9 (right) shows the distribution of the displaced pion $|d_{xy}|/\sigma_{d_{xy}}$ for signal and background events in the dimuon channel. Unlike candidates arising from HNL signals with a large lifetime, background candidates tend to be prompt, giving rise to a pion with limited displacement.

In the following, the preselection criteria are presented for the primary lepton, the displaced lepton, the displaced pion, and finally the $\ell_B \ell^\pm \pi^\mp$ candidate.

Primary lepton

The primary charged lepton, ℓ_B , must have $p_T > 2$ GeV and $|\eta| < 2$. Furthermore, additional criteria are imposed to improve the identification of this lepton. Muons must satisfy the so-called *soft* muon identification criteria. It means that they must be reconstructed with a high-quality track [126] that has at least one hit in the pixel detector and at least 5 hits in the strip detector, and that is matched to at least one segment in the muon detector. Moreover, as the primary




Figure 6.10: Efficiency and fake rate of the *loose* identification criteria for the displaced muon shown for the three benchmark signal hypotheses as a function of the displaced muon p_T (left) and the DV L_{xy} (right).

muon arises from the decay of a B meson and not from that of an HNL, its displacement is not large and the conditions $|d_z| < 20$ cm and $|d_{xy}| < 0.3$ cm are imposed. Electrons are identified using a discriminant based on a boosted decision tree architecture developed in the context of the work presented in Ref. [36].

Displaced lepton

The displaced charged lepton, $\ell^\pm$, is selected if it has $p_T > 2$ GeV and $|\eta| < 2$. Furthermore, it must satisfy $|d_z| > 0.0015$ cm, $|d_{xy}| > 0.001$ cm, and $|d_z|/\sigma_{d_z} > 1$, $|d_{xy}|/\sigma_{d_{xy}} > 1.5$, as this lepton arises from the decay of a long-lived heavy neutrino. If the lepton is a muon, it is required to be reconstructed with the particle-flow algorithm and to be either a *global* or *tracker* muon, defined in Section 2.4.4. These requirements are referred to as the *loose* muon identification criteria. The efficiency and fake rate of these criteria are shown in Fig. 6.10 for the three benchmark scenarios as a function of the displaced muon p_T and the DV L_{xy} . A loss of efficiency of about 10% is observed for $L_{xy} > 20$ cm. This effect is accounted for with a dedicated systematic uncertainty, as discussed in Section 6.8. The fake rate, defined as the rate of simulated particles misidentified as a muon, ranges between 1 and 15%. If the lepton is an electron, the same identification criteria as for the primary electron are applied.

Displaced pion

The displaced charged pion, $\pi^\mp$, must have $p_T > 1$ GeV and $|\eta| < 2$, as well as $|d_z| > 0.005$ cm, $|d_{xy}| > 0.005$ cm, and $|d_z|/\sigma_{d_z} > 1.5$, $|d_{xy}|/\sigma_{d_{xy}} > 5$. Moreover, to prevent the pion from being reconstructed with the same track as for the primary lepton, the condition $\Delta R(\ell_B, \pi^\mp) > 0.01$ is applied.

Signal candidate

Finally, the following preselection requirements are applied on the reconstructed $\ell_B \ell^\pm \pi^\mp$ signal candidates:

- $\text{charge}(\ell^\pm \pi^\mp) = 0$: the heavy neutrino is a neutral particle.
- $m(\ell_B \ell^\pm \pi^\mp) < 7 \text{ GeV}$: the invariant mass of the three objects in the final state is bounded from above by the mass of the B mesons.
- Trigger muon $p_T > 7 \text{ GeV}$ and $|\eta| < 1.5$: reinforcement of the B-parking trigger conditions. The effect of the triggers with p_T thresholds higher than 7 GeV is modelled with the trigger corrections described in Section 6.5.
- $p > 0.01$: p denotes the p -value associated with the χ^2 of the DV fit.
- $\cos \theta > 0.999$: the back-pointing angle θ is defined as the angle between the direction of the transverse momentum of the $\ell^\pm \pi^\mp$ system and the direction of flight of the HNL in the transverse plane. While the latter direction is unknown, as the decay vertex of the B meson is not reconstructed, it is well approximated by the direction of the vector joining the DV to the beamspot, given that the HNL is considered to travel a significantly larger distance than the B meson. As the heavy neutrino is considered to decay into visible particles, the back-pointing is compatible with 0 for signal events.
- $L_{xy}/\sigma_{L_{xy}} > 15$: the significance of the transverse displacement of the DV, defined as the ratio of L_{xy} with the associated error, $\sigma_{L_{xy}}$, is a key variable in the analysis as it is used to categorise the events, as discussed in the next section. The associated preselection cut is motivated by the fact that the heavy neutrino is long lived and decays from a vertex that is displaced with respect to the beamspot.
- $L_{xy} < 100 \text{ cm}$: the HNL is assumed to decay within the tracker volume.
- Vetoes on the known SM resonances listed in Section 6.7.2.

The preselection has an efficiency ranging from approximately 45% to 60% for the signal benchmark points, and an efficiency of about 5×10^{-6} for the background. The variables related to the kinematics and the displacement of the pion, as well as the back-pointing angle, θ , are among the variables that have the largest signal-to-background discriminating power. In an event, there is virtually always at most one reconstructed candidate passing the preselection conditions. In case there is more than one candidate, the one that has the largest $\cos \theta$ is kept for the analysis and the other ones are rejected.

6.4.2 Categories

After the preselection is applied, events are assigned to analysis categories with the aim of improving the sensitivity of the search by enhancing the signal significance of different HNL sce-

Table 6.5: List of the event categories.

Quantity	Label	Definition
$L_{xy}/\sigma_{L_{xy}}$	low $L_{xy}/\sigma_{L_{xy}}$	$L_{xy}/\sigma_{L_{xy}} < 50$
	medium $L_{xy}/\sigma_{L_{xy}}$	$50 < L_{xy}/\sigma_{L_{xy}} < 150$
	high $L_{xy}/\sigma_{L_{xy}}$	$L_{xy}/\sigma_{L_{xy}} > 150$
Relative lepton sign	OS	ℓ_B charge $\neq \ell$ charge
	SS	ℓ_B charge = ℓ charge
$m(\ell_B \ell^\pm \pi^\mp)$	low $m(\ell_B \ell^\pm \pi^\mp)$	$m(\ell_B \ell^\pm \pi^\mp) < 5.7$ GeV
	high $m(\ell_B \ell^\pm \pi^\mp)$	$m(\ell_B \ell^\pm \pi^\mp) > 5.7$ GeV
Flavour channel	dimuon	$\ell_B \ell = \mu\mu$
	mixed-flavour	$\ell_B \ell = (\mu e, e\mu)$

narios. In total, the events are divided into 24 mutually exclusive categories, listed in Table 6.5. These categories are based on

1. The significance of the transverse displacement, $L_{xy}/\sigma_{L_{xy}}$, of the DV with respect to the beamspot.
2. The relative sign of the primary and displaced leptons.
3. The $\ell_B \ell^\pm \pi^\mp$ invariant mass, $m(\ell_B \ell^\pm \pi^\mp)$.
4. The dimuon and mixed-flavour channels, introduced at the beginning of the chapter.

The categorisation in $L_{xy}/\sigma_{L_{xy}}$ is useful as it provides a way to enhance the sensitivity of the search to HNL signals with different lifetime hypotheses. Indeed, heavy neutrinos with long lifetimes are more likely to decay further away from the luminous region than heavy neutrinos with shorter lifetimes. As a consequence, the sensitivity to long-lived signals is enhanced in categories with large values of $L_{xy}/\sigma_{L_{xy}}$, while that to short-lived signals is enhanced in categories with small values of $L_{xy}/\sigma_{L_{xy}}$. Another advantage of the categorisation in $L_{xy}/\sigma_{L_{xy}}$ is related to the fact that the background, which mostly originates from prompt QCD processes, falls with larger values of $L_{xy}/\sigma_{L_{xy}}$, as shown in Fig. 6.11 (left). Therefore, the background population is smaller for large $L_{xy}/\sigma_{L_{xy}}$ values, increasing the signal significance in this region. The events are divided into three categories in $L_{xy}/\sigma_{L_{xy}}$: (i) $L_{xy}/\sigma_{L_{xy}} < 50$, (ii) $50 < L_{xy}/\sigma_{L_{xy}} < 150$, and (iii) $L_{xy}/\sigma_{L_{xy}} > 150$, referred to as the low, medium, and high $L_{xy}/\sigma_{L_{xy}}$ categories, respectively. Generally speaking, the sensitivity of the search is driven primarily by the high $L_{xy}/\sigma_{L_{xy}}$ categories for $1 < m_N < 3$ GeV, by the medium $L_{xy}/\sigma_{L_{xy}}$ categories for $3 < m_N < 5$ GeV, and by the low $L_{xy}/\sigma_{L_{xy}}$ categories for $m_N > 5$ GeV.

The events are also classified according to whether the primary and displaced leptons have the same sign (SS) or the opposite sign (OS), corresponding to lepton-number-violating and lepton-number-conserving decays, respectively. First of all, this categorisation is meant to enhance the sensitivity of the search to either Majorana or Dirac-like heavy neutrino hypotheses. For the case in which the HNL is a Majorana particle, the relative charge of the leptons is not constrained and both SS and OS lepton combinations are allowed. Instead, for the case in which the HNL is




Figure 6.11: Distribution of $L_{xy}/\sigma_{L_{xy}}$ (left) for data and the three signal benchmark hypotheses and distribution of the $\mu_B \mu^\pm \pi^\mp$ invariant mass (right) for (B_u , B_d , B_s) and B_c signals with $m_N = 4.5$ GeV and $c\tau_N = 1$ mm. The distributions are shown for events in the dimuon channel after the preselection and are normalised to unit area. The boundaries of the analysis categories are indicated with the orange lines. The last bin shows the overflow.

a Dirac-like particle, the leptons must be OS. Second of all, the background events are expected to populate the OS and SS categories differently. As the lepton number is conserved in SM decays, background events, arising for instance from B to D meson cascade decays, are expected to have predominantly OS leptons. For Majorana signals, the sensitivity is thus often driven by the SS categories.

Finally, the categorisation in the $\ell_B \ell^\pm \pi^\mp$ invariant mass is meant to distinguish between signals originating from the decay of a B_u , B_d , or B_s meson from those originating from the decay of a B_c meson. As discussed in Section 6.2.1, the B_c mesons give rise to an HNL signal only via the leptonic channel, implying that the $\ell_B \ell^\pm \pi^\mp$ invariant mass peaks at the mass of the meson, which is about 6.3 GeV. The other B meson species are lighter, with a mass of about 5.4 GeV. Therefore, signals arising from a B_u , B_d , or B_s meson can be selected with $m(\ell_B \ell^\pm \pi^\mp) < 5.7$ GeV, while those arising from a B_c meson can be selected with $m(\ell_B \ell^\pm \pi^\mp) > 5.7$ GeV. The two selections are referred to as the low $m(\ell_B \ell^\pm \pi^\mp)$ and high $m(\ell_B \ell^\pm \pi^\mp)$ categories, respectively, and are indicated in Fig. 6.11 (right), which shows the distribution of $m(\ell_B \ell^\pm \pi^\mp)$ for (B_u , B_d , B_s) and B_c signal events simulated with $m_N = 4.5$ GeV and $c\tau_N = 1$ mm.

6.4.3 Selection

A strength of the search is that it probes heavy neutrinos over a fine grid of mass and lifetime hypotheses: over a hundred masses are tested in the range $1 < m_N < 6$ GeV and, for each of them, about 20 lifetime scenarios are considered. This strength comes with the complexity that each of these signal hypotheses can give rise to final states with significantly different signatures. This behaviour is illustrated in Fig. 6.7, where distributions obtained for different signal hypotheses are compared. Taking as an example the $m(\ell_B \ell^\pm \pi^\mp)$ distribution, we see that

the selection condition $m(\ell_B \ell^\pm \pi^\mp) < 3$ GeV would be a good handle to select signal events with $m_N \lesssim 3$ GeV, while it would discard signal events with $m_N \gtrsim 3$ GeV. This example shows that the use of a unique set of selection conditions is not well suited for the purposes of this search. In fact, one of the main challenges of the analysis was the design of an event selection algorithm optimal for all the signal m_N and $c\tau_N$ hypotheses.

Parametric neutral network

The approach found to optimise the selection performance for all the m_N hypotheses is based on a multivariate discriminator that incorporates a functional dependence on m_N , known as a parametric neural network (pNN) [115]. This type of neural network (NN) takes as input one or several discrete parameters, denoted as $\vec{\theta}$, in addition to a set features, $\vec{x}$, which are the variables used to discriminate between signal and background events. When evaluated on a given event, the pNN returns a score that is a function of the full variable set, $f(\vec{\theta}, \vec{x})$. In the following, only one discrete parameter, m_N , is chosen, i.e. $\vec{\theta} = \{m_N\}$. The advantages of the pNN are that (i) a single training can be performed with different values of the parameter m_N instead of having separate NNs trained for each of these values, that (ii) the performance obtained for each of the values of m_N is as good as for an NN trained for that mass only, and that (iii) the pNN can well interpolate between the different values of m_N , implying that it can be trained on a handful of values of m_N instead of the ~ 100 values of m_N probed in the search.

The strategy established in this analysis consists in training pNNs, with m_N as a parameter, in the different analysis categories, listed in Table 6.5. The value of the m_N parameter, hereafter referred to as mass keys, used for the training are $m_N = 1.0, 1.5, 2.0, 3.0$, and 4.5 GeV for the low $m(\ell_B \ell^\pm \pi^\mp)$ categories and $m_N = 3.0$ and 4.5 GeV in the high $m(\ell_B \ell^\pm \pi^\mp)$ categories. For the latter, the parameters $m_N = 1.0, 1.5$, and 2.0 GeV are not used, as signals from B_c mesons were not simulated for these masses, as discussed in Section 6.2. The variable $c\tau_N$ is not used as a parameter of the pNN, as it was found that training the pNN in the different $L_{xy}/\sigma_{L_{xy}}$ categories, using signal events with mixed values of $c\tau_N$, maintained a good performance for all the $c\tau_N$ hypotheses, as further discussed below.

In each of the analysis categories, the pNN is trained on signal and background events passing the preselection criteria described in Section 6.4.1. Signal events simulated for each of the mass keys and with different values of $c\tau_N$ in the range $0.1 < c\tau_N < 1000$ mm are aggregated into a common signal sample used for the pNN training. For the background sample, events are taken out of a small part of the B-parking data set, corresponding to less than one thousandth of the available data events. In fact, exclusion limits from early HNL searches exclude the possibility of the presence in such a small sample of statistically significant signal events that would bias the training. Because some of the input parameters, $\vec{x}$, depend on m_N , the background training sample cannot be chosen as a random sample of data, as proposed in Ref. [115]. Instead, for each mass key, the background events are selected for the training if the corresponding $\ell^\pm \pi^\mp$ invariant mass lies within a window of size $\pm 10\sigma$ around the corresponding value of m_N , where σ is the signal mass resolution discussed in Section 6.6.6.

After the pNN is trained to learn the differences between signal and background events, it can be used in data to enhance the signal significance with respect to the background. For each event, the pNN is evaluated for a given mass parameter, m_N , and returns a score, $f(\vec{\theta}, \vec{x})$. The pNN score assigned to an event can be interpreted as the probability for this event to originate from a signal process with mass m_N . Signal-like events are selected in the analysis by requiring the pNN score to be greater than 0.99. This threshold was obtained from a dedicated study in which the performance obtained using different cuts on the pNN score was evaluated based on the median expected limit. It was found to perform well for all the event categories.

Input features

The pNN is trained using a set of features, $\vec{x}$, that provide a good discrimination between signal and background events. The distribution of some of these features is shown in Fig. 6.12 for background and signal events with $m_N = 2$ GeV and $c\tau_N = 100$ mm, in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel. These distributions are discussed below. The complete set of features is the following:

1. Transverse momenta: $p_T(\ell_B)$, $p_T(\ell^\pm)$, $p_T(\pi^\mp)$.

These variables are useful as the p_T distribution of the background tends to be softer than that of the signal, as can be seen in Fig. 6.12 (*upper left*). The discrimination power of these quantities depends on the signal mass, m_N . For instance, the distribution of $p_T(\ell^\pm)$ and $p_T(\pi^\mp)$ becomes harder with larger values of m_N , unlike that of $p_T(\ell_B)$ that becomes softer.

2. Invariant masses: $m(\ell_B \pi^\mp)$, $m(\ell_B \ell^\pm)$, $m(\ell_B \ell^\pm \pi^\mp)$.

As discussed in Section 6.2.3, the invariant mass spectra strongly depend on the value of m_N . For large masses, the B meson decays leptonically and its invariant mass can be fully reconstructed, providing a handle to further reject background events. For smaller masses, as shown in Fig. 6.12 (*upper right*), the decays are mostly semileptonic and the $m(\ell_B \ell^\pm \pi^\mp)$ distribution does not peak at the mass of the B meson, as the accompanying hadronic system, X, is not reconstructed.

3. Track separation in the η - ϕ space: $\Delta R(\ell_B, \ell^\pm)$, $\Delta R(\ell_B, \pi^\mp)$.

These variables offer a good separation power between signal and background processes, as the final state particles in a background event tend to be less aligned than these in a signal event, yielding larger values of ΔR . The variable $\Delta R(\ell^\pm, \pi^\mp)$ is not added to the feature set, as otherwise the HNL mass could be reconstructed during the training as the transverse momenta $p_T(\ell^\pm)$ and $p_T(\pi^\mp)$ are also input features. Thorough checks have shown that the pNN does not learn the mass m_N during the training.

4. Displaced vertex properties: $\cos \theta$, fit p -value.

The addition of the cosine of the back-pointing angle, θ , whose distribution is shown in Fig. 6.12 (*lower left*), provides additional sensitivity to the $c\tau_N$ hypothesis. Moreover, the

distribution of this variable typically peaks at 1 more steeply for signal than for background events. The addition of the fit p -value is useful as background events tend to give rise to larger values than signal events.

5. Displacement-related quantities: $L_{xy}/\sigma_{L_{xy}}$ and $|d_{xy}|/\sigma_{d_{xy}}$ of the pion.

These quantities depend on the lifetime and allow the pNN to have a good performance for the different $c\tau_N$ hypotheses. The distribution of the pion $|d_{xy}|/\sigma_{d_{xy}}$ is shown in Fig. 6.12 (lower right) and tends to have larger values for the long-lived signal than for the background.

6. Track-related information: number of layers of the pixel and strip detectors used to reconstruct the track of the final state particles.

Information on the tracks is provided to the pNN as it gives another handle to discriminate between the long-lived signal and the background. Particles originating from the decay of an HNL with large $c\tau_N$ typically traverse fewer tracker layers than particles originating from prompt background processes.

7. Lepton isolation, defined in a cone of size $\Delta R = 0.3$ around the lepton $\vec{p}_T$.

The PF isolation, defined in Section 2.4.4, of the leptons is used as a training feature as leptons originating from a signal process tend to be more isolated than leptons originating from a background process.

Training

The architecture of the pNN is that of a fully connected NN with three layers:

1. Input layer: contains as many nodes as the number of features, plus one extra node for the parameter, m_N . The method used in the analysis hence consists in concatenating the parameter to the set of input features.
2. Hidden layer: contains 64 nodes and is activated by the ReLU function, defined as $\text{ReLU}(x) = \max(0, x)$, with $x \in \mathbb{R}$.
3. Output layer: returns the pNN score, normalised to one through the sigmoid activation function, defined as $\text{sigmoid}(x) = 1/(1 + \exp x)$, with $x \in \mathbb{R}$.

The training is performed using about 50×10^3 to 100×10^3 signal-plus-background events. The samples contain roughly the same number of events associated with each of the mass keys. The training is executed with the TENSORFLOW package [127], using the KERAS interface [128] and the ADAM optimiser [129]. The batch size is set to 32, and the initial learning rate to 0.01, but can be further reduced during the training if the accuracy does not improve over some epochs. The pNN is typically trained over 100 epochs. However, in case the performance reaches a plateau, the training is stopped earlier, to avoid the possibility of over-training the network.




Figure 6.12: The distributions of the pion p_T (upper left), $m(\mu_B \mu^\pm \pi^\mp)$ (upper right), $\cos \theta$ (lower left), and pion $|d_{xy}|/\sigma_{d_{xy}}$ (lower right) are shown for data and for a signal hypothesis of $m_N = 2$ GeV and $c\tau_N = 100$ mm. The data are selected in the mass window of size $\pm 10\sigma$ around $m_N = 2$ GeV and correspond to an integrated luminosity of 5.2 fb^{-1} . The distributions, which are normalised to unit area, are shown in the category with high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ in the dimuon channel.

Performance

After the pNN is trained on events in a given search category, the performance is obtained by evaluating it on different signal and background events than those used for the training. The comparison of signal and background distributions of the pNN score is shown in Fig. 6.13 (left), for signal hypotheses with $m_N = 1.0$ GeV and $c\tau_N = 10.0, 100.0$, and 1000.0 mm and events in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category of the dimuon channel. We observe that the signal distributions peak at 1, while the background one peaks at 0, which shows the good signal-to-background separation power of the pNN. The condition that the pNN score be greater than 0.99 therefore greatly enhances the signal significance with respect to the background. This behaviour is confirmed in Fig. 6.13 (right), which shows the signal efficiency against the




Figure 6.13: Shape of the distribution of the pNN score (*left*) and ROC curve (*right*) for signal hypotheses with $m_N = 1.0$ GeV and $c\tau_N = 10.0, 100.0$, and 1000.0 mm for the pNN trained in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel.

background efficiency with a scan on the pNN score (ROC curve). For the aforementioned signal hypotheses and category, a cut on the score at 0.99 preserves approximately 50% of the signal, while it rejects about 99.9% of the background. Generally speaking, the efficiency of the cut on the pNN score ranges across the different event categories and signal hypotheses from approximately 30 to 70% for the signal and from approximately 0.01 to 1% for the background. The efficiencies are obtained after the preselection is applied.

The performance of the pNN-based selection algorithm was evaluated against the performance of a cut-based selection. The latter selection consists of cuts applied sequentially on $m(\ell_B \ell^\pm \pi^\mp)$, $p_T(\pi^\mp)$, $L_{xy}/\sigma_{L_{xy}}$, and $\cos \theta$, and optimised against the signal significance for $m_N = 1.0, 3.0$, and 4.5 GeV, corresponding to the semileptonic, mixed, and leptonic regimes, respectively, as shown in Fig. 6.4 (*left*). It was found that the performance of the pNN outperformed that of the cut-based approach for the three mass regimes and the different $c\tau_N$ hypotheses: for the same signal efficiency, the background efficiency obtained with the pNN was lowered by up to one order of magnitude compared to that obtained with the cut-based strategy.

Although the pNN is not trained using $c\tau_N$ as a parameter, the training strategy employed in the analysis provides a good performance for the different $c\tau_N$ hypotheses. This is verified in Fig. 6.14 (*left*), which shows the area under the ROC curve (AUC) for the different values of m_N used as mass keys during the training, and for different $c\tau_N$ hypotheses. The events are selected in the medium $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel. It is shown that the performance of the pNN is good, i.e. AUC values are close to 1, over a wide range of $c\tau_N$ values, with a minor degradation for signals with smaller $c\tau_N$.

To validate the performance of the pNN against that of a regular NN, we compare the AUC values obtained with the two approaches. The results of this study are shown in Fig. 6.14 (*right*) for events in the medium $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel.




Figure 6.14: (Left) Performance of the pNN assessed with the AUC as a function of m_N for events in the $50 < L_{xy}/\sigma_{L_{xy}} < 150$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel. The AUC at unity corresponds to a perfect separation between signal and background. The different coloured curves show different $c\tau_N$ hypotheses. (Right) Comparison of the AUC for a pNN trained on masses $m_N = 1.0, 1.5, 2.0, 3.0$, and 4.5 GeV (blue) and for a NN trained on mass $m_N = 2.0 \text{ GeV}$ (red), for events in the $50 < L_{xy}/\sigma_{L_{xy}} < 150$, OS and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel. The points have $c\tau_N = 10 \text{ mm}$. The full circles show mass points on which the pNN and NN were trained on, while the open circles show mass points that were not trained on.

The first set of AUC values is obtained using a pNN trained on masses $m_N = 1.0, 1.5, 2.0, 3.0$, and 4.5 GeV , while the second set of AUC values is obtained using an NN with the same architecture and training conditions as the pNN, but trained on mass $m_N = 2.0 \text{ GeV}$ only. The signal points have $c\tau_N = 10 \text{ mm}$. For the common mass on which the two networks are trained, $m_N = 2.0 \text{ GeV}$, we observe that the pNN and the NN yield the same AUC score, which shows that the pNN does not lose performance compared to an NN trained on a single mass. Moreover, we see that the performance of the pNN remains approximately constant for all masses, unlike the regular NN, which loses performance when evaluated on masses away from 2.0 GeV . Finally, the pNN is also evaluated for masses that were not used during the training. We observe that the AUC values obtained for these masses are high, showing that the pNN is able to well interpolate between masses it was trained on.

6.5 Monte Carlo corrections

Simulated signal events have a key role in the analysis, notably as they are used to obtain the signal efficiency needed to compute the expected signal yield, as described in Section 6.6. It is thus important to make sure that the simulation can well reproduce the data without the introduction of unwanted biases, and to compute and apply corrections when necessary.

First of all, the modelling of the displacement-related variables used in the pNN selection, $L_{xy}/\sigma_{L_{xy}}$ and $|d_{xy}|/\sigma_{d_{xy}}(\pi^\mp)$, was checked in MC using the well-defined SM decay $J/\psi \rightarrow \mu\mu$. As the background is taken from data and the signal from simulation during the pNN training,

this study is used to ensure that the network learns the genuine differences between signal and background events rather than the possible mismodelling of the simulated events compared to data. It was found that both $L_{xy}/\sigma_{L_{xy}}$ and $|d_{xy}|/\sigma_{d_{xy}}(\pi^\mp)$ were well modelled in the simulation; thus no corrections were applied.

Second, a series of weights was derived to correct for some inaccuracies in the MC samples. These weights are applied on each simulated event and impact the expected signal yield, rather than the signal shape. These corrections are

1. Corrections for the pileup distribution. These corrections are minor, as the pileup profile injected during the simulation of the events is close to that in data, as mentioned in Section 6.2.2.
2. Corrections for the efficiency of the *soft* and *loose* muon identification criteria, defined in Section 6.4.1. These corrections were derived within the muon group at CMS and are typically small.
3. Corrections for the trigger efficiency. These corrections are by far the largest ones and are further discussed below.

As described in Section 6.1, the B-parking sample was accumulated using a series of ten muon trigger lines with different thresholds on the muon p_T and $|d_{xy}|/\sigma_{d_{xy}}$. The data are collected in such a way that the trigger thresholds are adjusted according to the instantaneous luminosity. As a consequence, not all the trigger lines were active at the same time, i.e. the trigger prescales were not constant, giving rise to a complex trigger structure. When simulating the signal samples, following the description given in Section 6.2, the dependence of the trigger thresholds on the instantaneous luminosity was not accounted for. This effect is clearly visible in Fig. 6.15 (left), which shows that the shape of the distribution of the p_T of the triggering μ_B in MC does not overlap with that in data. The MC events are simulated for soft QCD processes with muons. The distributions are shown for events in the dimuon channel reconstructed and preselected as described in Sections 6.3 and 6.4.1 for signal, except that the charge of the $\mu^\pm\pi^\mp$ system is non-zero. Therefore, corrections must be computed such that the effect of the complex trigger composition used to collect the B-parking sample is reproduced in the simulated signal samples.

For an event in the mixed-flavour channel, in which there is exactly one muon in the final state, the correction, w , is applied as an event weight obtained as

$$w = \frac{\varepsilon^{\text{data}}}{\varepsilon^{\text{MC}}}, \quad (6.2)$$

where $\varepsilon^{\text{data}}$ and ε^{MC} are the trigger efficiencies in data and simulation, respectively. These efficiencies are derived in bins of the muon p_T and $|d_{xy}|/\sigma_{d_{xy}}$ using the method described below. For an event in the dimuon channel, either one or both muons may fire a B-parking trigger, and




Figure 6.15: Distribution of the p_T of the triggering μ_B shown without (left) and with (right) the application of the corrections for the trigger efficiency. The distributions are compared between data and simulated soft QCD events in the dimuon channel with a non-zero charge of the $\mu^\pm\pi^\mp$ system.

the correction is obtained as

$$w = \begin{cases} \frac{\varepsilon_B^{\text{data}}}{\varepsilon_B^{\text{MC}}} \cdot \frac{1-\varepsilon^{\text{data}}}{1-\varepsilon^{\text{MC}}} & \text{if } \mu_B \text{ triggers and } \mu^\pm \text{ does not trigger,} \\ \frac{1-\varepsilon_B^{\text{data}}}{1-\varepsilon_B^{\text{MC}}} \cdot \frac{\varepsilon^{\text{data}}}{\varepsilon^{\text{MC}}} & \text{if } \mu_B \text{ does not trigger and } \mu^\pm \text{ triggers,} \\ \frac{\varepsilon_B^{\text{data}}}{\varepsilon_B^{\text{MC}}} \cdot \frac{\varepsilon^{\text{data}}}{\varepsilon^{\text{MC}}} & \text{if both } \mu_B \text{ and } \mu^\pm \text{ trigger,} \end{cases} \quad (6.3)$$

where the symbol ε_B denotes the efficiencies associated with the primary muon, μ_B , and ε denotes the efficiencies associated with the displaced muon, $\mu^\pm$.

The trigger efficiencies are computed in both data and MC using the *tag-and-probe* method [130] with $J/\psi \rightarrow \mu\mu$ decays. Events are selected with the requirements that the two muons have $p_T > 1$ GeV and $|\eta| < 2.3$, and that they have opposite charge. The muon tracks are fitted to a common vertex such that $\cos\theta > 0.9$. Moreover, the dimuon invariant mass must agree within 200 MeV with the mass of the J/ψ . For data, the entire luminosity of the B-parking sample is used, while a dedicated sample was simulated for the $B \rightarrow J/\psi(\mu\mu)K^*$ process for MC. One muon, referred to as the *tag* muon, is further required to be identified with the *soft* criteria introduced in Section 6.4.1 and to fire one or both of the B-parking triggers with thresholds (i) $p_T = 9$ GeV and $|d_{xy}|/\sigma_{d_{xy}} = 6$, or (ii) $p_T = 12$ GeV and $|d_{xy}|/\sigma_{d_{xy}} = 6$. These two trigger lines are used as at least one of them was active during the collection of the B-parking sample. The other muon, referred to as the *probe* muon, is left unconstrained, such that the efficiency can be measured in an unbiased way.

The efficiencies are obtained in bins of the probe muon p_T and $|d_{xy}|/\sigma_{d_{xy}}$, with boundaries chosen such that they correspond to the different HLT trigger thresholds. In each p_T and $|d_{xy}|/\sigma_{d_{xy}}$ bin, for both data and MC, the efficiency is computed as the number of probe muons firing any of the B-parking triggers divided by the total number of probe muons. These numbers




Figure 6.16: Ratio of the trigger efficiencies in data and MC in bins of the probe muon p_T and $|d_{xy}|/\sigma_{d_{xy}}$. The uncertainty corresponds to the propagated uncertainty in the number of events returned by the fits to the dimuon mass spectra in data and simulation.

are obtained with an unbinned extended maximum likelihood fit to the dimuon invariant mass distributions.

The ratio of the efficiency in data to that in MC is shown in Fig. 6.16 for each p_T and $|d_{xy}|/\sigma_{d_{xy}}$ bin. As expected, the value of the ratio is small in bins corresponding to trigger lines with low thresholds and it gets close to 1 after all the triggers have reached the efficiency plateau.

The trigger corrections are validated with the comparison of the distributions of $p_T(\mu_B)$ for data and MC events shown in Fig. 6.15 (right). While the corrections were computed using the $J/\psi \rightarrow \mu\mu$ sample, the closure is studied using a different simulated sample, corresponding to QCD events enriched with muons, to check the robustness of the corrections. With the trigger corrections applied, the distribution for simulated events is in good agreement with that for data, which confirms the sanity of the corrections. However, a shortcoming of the method is that the efficiencies are obtained using the decay of a J/ψ that is short lived compared to the HNL signals probed in the search. As a consequence, the bins at high displacement are not well populated, which is accounted for with a dedicated systematic uncertainty, as discussed in Section 6.8.

6.6 Signal normalisation and parametrisation

To extract the signal from the data and to obtain exclusion limits on the parameters describing the phenomenology of the search, the number of expected signal events must be precisely computed. The general procedure to normalise the signal yield is presented in Section 6.6.1. In particular, the methods to reweight the number of expected signal events to arbitrary mixing sce-

narios, to the Dirac-like case, and to arbitrary lifetime hypotheses are described in Sections 6.6.2, 6.6.3, and 6.6.4, respectively. As discussed in Section 6.6.1, a key component to normalise the signal events is the effective cross section for the production of $B^\pm$ in the B-parking data sample, whose measurement is presented in Section 6.6.5. Finally, the signal is extracted from the data according to the shape of the distribution of the invariant mass of the displaced objects, following the parametrisation method described in Section 6.6.6.

6.6.1 Signal yield normalisation

In the lepton flavour channel $\ell_B \ell = (\mu\mu, e\mu, \mu e)$, the number of expected signal events, N_{sig} , for a given signal hypothesis described by m_N and $c\tau_N$, and for the mixing scenario specified by $\vec{r} = (r_e, r_\mu, r_\tau)$, is computed as

$$N_{\text{sig}}(\ell_B \ell, m_N, c\tau_N, \vec{r}) = \frac{\sigma_{B^\pm}^{\text{eff}}}{f_u} \times \mathcal{L} \times \sum_q F_N^q(\ell_B \ell, m_N, c\tau_N, \vec{r}) \times \epsilon_{\text{sig}}^q(\ell_B \ell, m_N, c\tau_N), \quad (6.4)$$

where

- $\sigma_{B^\pm}^{\text{eff}}$ is the effective total cross section for the production of $B^\pm$ mesons in the kinematic phase space covered by the B-parking trigger conditions. The measurement of $\sigma_{B^\pm}^{\text{eff}}$ is presented in Section 6.6.5.
- f_u is the fragmentation fraction of the B_u meson, given in Table 6.3.
- $\mathcal{L}$ is the integrated luminosity for the B-parking data set equal to 41.6 fb^{-1} .
- q is the quark flavour index and runs over (u, d, s, c).
- F_N^q is the signal factor that accounts for the HNL production rate associated with a signal coming from the decay of a meson B_q and for the HNL decay branching fraction, as explained below.
- ϵ_{sig}^q is the associated signal efficiency obtained from simulation. It comprises the acceptance efficiency, applied with the event filters described in Section 6.2.2; the reconstruction efficiency, described in Section 6.3; the selection efficiency, described in Section 6.4; as well as the effect from the MC corrections, described in Section 6.5. The efficiencies are obtained separately for signals originating from B_u , B_d , B_s , or B_c decays, as anticipated in Section 6.2.2. Moreover, the effect of the efficiency of the matching of reconstructed to generator-level candidates, discussed in Section 6.3, is corrected for.

The signal factor, F_N^q , is defined as

$$F_N^q(\ell_B \ell, m_N, c\tau_N, \vec{r}) = \underbrace{\sum_{X_b} f_q \frac{\Gamma(B_q \rightarrow \ell_B N X_b)}{\Gamma(B_q)}}_{\text{HNL production}} \times \underbrace{\frac{\Gamma(N \rightarrow \ell \pi)}{\Gamma_e(N) + \Gamma_\mu(N) + \Gamma_\tau(N)}}_{\text{HNL decay}} r_\ell, \quad (6.5)$$

where the first term describes the HNL production and the second term describes the HNL decay. In the HNL production term, the sum runs over the possible accompanying mesons, X_b ,

associated with the decay of the meson B_q and listed in Table 6.2, and f_q is the fragmentation fraction of the meson B_q reported in Table 6.3. The numerator of the ratio, $\Gamma(B_q \rightarrow \ell_B N X_b)$, is the decay rate, described in Section 6.2.1, for the production of the HNL along with the primary lepton, ℓ_B , and the denominator of the ratio, $\Gamma(B_q)$, is the SM B_q decay width [3]. In the HNL decay term, the factor $\Gamma(N \rightarrow \ell\pi)$ is the partial decay width for an HNL decay into $\ell\pi$, and $\Gamma_\alpha(N)$ is the total decay width for an HNL whose decay involves a charged or neutral lepton of flavour $\alpha = (e, \mu, \tau)$. The HNL decay widths are described in Section 6.2.1. The dependence in $|V_N|^2$ and $\vec{r}$ of the signal factor can be made explicit using the reduced widths, $\tilde{\Gamma}$, defined as

$$\tilde{\Gamma}(B_q \rightarrow \ell_B N X_b) = \frac{1}{|V_{\ell_B N}|^2} \Gamma(B_q \rightarrow \ell_B N X_b), \quad (6.6)$$

$$\tilde{\Gamma}(N \rightarrow \ell\pi) = \frac{1}{|V_{\ell N}|^2} \Gamma(N \rightarrow \ell\pi), \quad (6.7)$$

$$\tilde{\Gamma}_\alpha(N) = \frac{1}{|V_{\alpha N}|^2} \Gamma_\alpha(N). \quad (6.8)$$

Using $|V_{\alpha N}|^2 = r_\alpha |V_N|^2$, the signal factor, F_N^q , can hence be expressed as

$$F_N^q(\ell_B \ell, m_N, c\tau_N, \vec{r}) = \underbrace{\sum_{X_b} f_q \frac{\tilde{\Gamma}(B_q \rightarrow \ell_B N X_b)}{\Gamma(B_q)} r_{\ell_B} |V_N|^2}_{\text{HNL production}} \times \underbrace{\frac{\tilde{\Gamma}(N \rightarrow \ell\pi)}{r_e \tilde{\Gamma}_e(N) + r_\mu \tilde{\Gamma}_\mu(N) + r_\tau \tilde{\Gamma}_\tau(N)}}_{\text{HNL decay}} r_\ell. \quad (6.9)$$

This expression will be useful to understand the reweighting to different mixing scenarios, discussed in the next section.

6.6.2 Reweighting to arbitrary mixing scenarios

As described in Section 6.2.2, each of the signal MC samples is simulated for

1. A fixed m_N hypothesis.
2. A fixed $c\tau_N$ hypothesis.
3. The Majorana scenario.
4. The mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ for events in the dimuon channel and the mixing scenario $(r_e, r_\mu, r_\tau) = (1/2, 1/2, 0)$ for events in the mixed-flavour channel. These scenarios are referred to as the *nominal* mixing scenario, $(r_e^0, r_\mu^0, r_\tau^0)$.

A strength of the analysis is that the results are interpreted for more than 60 mixing scenarios, while most of the searches so far have constrained merely the 3 exclusive benchmark scenarios, in which the HNL mixes only with one neutrino flavour. Instead of simulating signal events for each of these mixing scenarios to estimate the signal yields, weights are applied on the events generated with the nominal conditions. The derivation of the weight to obtain the expected number of signal events for the *target* mixing scenario, $(r_e^{\text{target}}, r_\mu^{\text{target}}, r_\tau^{\text{target}})$, is presented in the following.

First of all, one must take into account that, for a fixed m_N and $c\tau_N$, a change of the mixing scenario implies a change in $|V_N|^2$. Indeed, because $c\tau_N$ is fixed, i.e.

$$\frac{1}{c\tau_N} = \sum_{\alpha} \Gamma_{\alpha} = |V_N|^2 \sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha} = \text{const}, \quad (6.10)$$

where α runs over (e, μ, τ) , we can obtain $|V_N|^2$ associated with the target mixing scenario, $|V_N|^{2, \text{target}}$, from the one associated with the nominal mixing scenario, $|V_N|^{2, 0}$, as

$$|V_N|^{2, \text{target}} = |V_N|^{2, 0} \times \frac{\sum_{\alpha} r_{\alpha}^0 \tilde{\Gamma}_{\alpha}}{\sum_{\alpha} r_{\alpha}^{\text{target}} \tilde{\Gamma}_{\alpha}}. \quad (6.11)$$

Then, based on Eq. (6.9), the number of expected signal events, $N_{\text{sig}}^{\text{target}}$, associated with the target mixing scenario can be obtained in each analysis category from the number of events, N_{sig}^0 , associated with the nominal scenario, as

$$N_{\text{sig}}^{\text{target}} = N_{\text{sig}}^0 \times \frac{|V_N|^{2, \text{target}}}{|V_N|^{2, 0}} \times \frac{r_{\ell_B}^{\text{target}}}{r_{\ell_B}^0} \times \frac{r_{\ell}^{\text{target}}}{r_{\ell}^0} \times \frac{\sum_{\alpha} r_{\alpha}^0 \tilde{\Gamma}_{\alpha}}{\sum_{\alpha} r_{\alpha}^{\text{target}} \tilde{\Gamma}_{\alpha}}. \quad (6.12)$$

Substituting Eq. (6.11) in Eq. (6.12), we find

$$N_{\text{sig}}^{\text{target}} = N_{\text{sig}}^0 \times \frac{r_{\ell_B}^{\text{target}}}{r_{\ell_B}^0} \times \frac{r_{\ell}^{\text{target}}}{r_{\ell}^0} \times \left(\frac{\sum_{\alpha} r_{\alpha}^0 \tilde{\Gamma}_{\alpha}}{\sum_{\alpha} r_{\alpha}^{\text{target}} \tilde{\Gamma}_{\alpha}} \right)^2. \quad (6.13)$$

6.6.3 Reweighting to the Dirac-like scenario

As discussed above, the signal samples are simulated for the Majorana scenario, implying that the relative sign of the two leptons, ℓ_B and $\ell^{\pm}$, is left unconstrained. We explain in the following that these samples can be used to obtain the number of expected signal events for the Dirac-like scenario.

If the heavy neutrino is a Dirac-like particle, its decay width is half of that for a Majorana particle [95]. Therefore, using Eq. (6.10), we can write

$$\begin{aligned} |V_N|^{2, \text{Maj.}} \sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha}^{\text{Maj.}} &= |V_N|^{2, \text{Dir.}} \sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha}^{\text{Dir.}} \\ &= |V_N|^{2, \text{Dir.}} \frac{1}{2} \sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha}^{\text{Maj.}}, \end{aligned} \quad (6.14)$$

and deduce that

$$|V_N|^{2, \text{Dir.}} = 2|V_N|^{2, \text{Maj.}}, \quad (6.15)$$

where the subscripts “Maj.” and “Dir.” refer to the Majorana and Dirac-like scenarios, respectively. Moreover, lepton-number-violating decays are not allowed if the HNL is a Dirac-like particle, which implies that the signal yields in the SS categories are set to zero,

$$N_{\text{sig}}^{\text{Dir.}}(\text{SS}) = 0. \quad (6.16)$$

Finally, for the signal yield in the OS categories, the number of signal events associated with a Dirac-like particle can be obtained from those associated with a Majorana particle as

$$\begin{aligned}
 N_{\text{sig}}^{\text{Dir.}}(\text{OS}) &= N_{\text{sig}}^{\text{Maj.}}(\text{OS}) \times \frac{|V_N|^2, \text{Dir.}}{|V_N|^2, \text{Maj.}} \times \frac{\tilde{\Gamma}^{\text{Dir.}}(N \rightarrow \ell\pi)}{\tilde{\Gamma}^{\text{Maj.}}(N \rightarrow \ell\pi)} \times \frac{\sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha}^{\text{Maj.}}}{\sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha}^{\text{Dir.}}} \\
 &= N_{\text{sig}}^{\text{Maj.}}(\text{OS}) \times \frac{2|V_N|^2, \text{Maj.}}{|V_N|^2, \text{Maj.}} \times \frac{1/2 \tilde{\Gamma}^{\text{Maj.}}(N \rightarrow \ell\pi)}{\tilde{\Gamma}^{\text{Maj.}}(N \rightarrow \ell\pi)} \times \frac{\sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha}^{\text{Maj.}}}{1/2 \sum_{\alpha} r_{\alpha} \tilde{\Gamma}_{\alpha}^{\text{Maj.}}} \\
 &= 2N_{\text{sig}}^{\text{Maj.}}(\text{OS}).
 \end{aligned} \tag{6.17}$$

6.6.4 Reweighting to arbitrary lifetimes

As discussed in Section 6.2.2, for each m_N hypothesis, signal samples are generated for a handful of $c\tau_N$ hypotheses. To assess the sensitivity of the search for a finer grid of $c\tau_N$ points, a weight is applied on simulated events to reproduce the underlying ct distribution associated with the target $c\tau_N$ value. This procedure is referred to as the $c\tau_N$ -reweighting.

The $c\tau_N$ -reweighting is based on the fact that the lifetime, t , of the HNL in the laboratory frame distributes as an exponential function with parameter $1/c\tau_N$. Therefore, events simulated with $c\tau_N^{(1)}$ can be obtained from those simulated with $c\tau_N^{(0)}$ using the weight

$$w\left(ct, c\tau_N^{(0)} \rightarrow c\tau_N^{(1)}\right) = \frac{\frac{1}{c\tau_N^{(1)}} \exp\left(-\frac{ct}{c\tau_N^{(1)}}\right)}{\frac{1}{c\tau_N^{(0)}} \exp\left(-\frac{ct}{c\tau_N^{(0)}}\right)}. \tag{6.18}$$

The value of ct is calculated at the generator level using $ct = L_{\text{xyz}}/\beta_N\gamma_N$, where L_{xyz} is the displacement of the HNL and $\beta_N\gamma_N$ is its boost factor to the laboratory frame.

In principle, any value of $c\tau_N^{(1)}$ can be obtained from a sample generated for the hypothesis $c\tau_N^{(0)}$. However, in practice, there must be enough events in the ct distribution associated with the parameter $1/c\tau_N^{(0)}$ for a proper $c\tau_N$ -reweighting, especially when a small value of $c\tau_N^{(0)}$ is used to reweight events to significantly larger values of $c\tau_N^{(1)}$. Therefore, instead of using a unique sample with hypothesis $c\tau_N^{(0)}$ to perform the reweighting, all the samples $\{i\}$ generated for a given m_N and for $c\tau_N^{(i)}$ are exploited. The expression in Eq. (6.18) can be generalised to

$$w\left(ct, \{c\tau_N^{(i)}\} \rightarrow c\tau_N^{(1)}\right) = \frac{\frac{1}{c\tau_N^{(1)}} \exp\left(-\frac{ct}{c\tau_N^{(1)}}\right)}{\sum_i \frac{N_{\text{gen}}^{(i)}}{N_{\text{gen}}^{\text{tot}}} \frac{1}{c\tau_N^{(i)}} \exp\left(-\frac{ct}{c\tau_N^{(i)}}\right)}, \tag{6.19}$$

where $N_{\text{gen}}^{(i)}$ is the number of generated events for the sample with hypothesis $c\tau_N^{(i)}$ and $N_{\text{gen}}^{\text{tot}}$ is the total number of generated events for all the $c\tau_N$ hypotheses.

6.6.5 Measurement of the effective cross section for $B^{\pm}$ production

The total effective cross section, $\sigma_{B^{\pm}}^{\text{eff}}$ in Eq. (6.4), is defined such that $\sigma_{B^{\pm}}^{\text{eff}} \times \mathcal{L}$ represents the number of proton-proton collisions giving rise to a B^+ or B^- meson in the B-parking data sample.

It is obtained from the extrapolation of the measurement of the cross section for the control SM process $B^\pm \rightarrow J/\psi(\rightarrow \mu^+\mu^-)K^\pm$, performed in the fiducial phase space region imposed by the B-parking trigger requirements.

The process $B^\pm \rightarrow J/\psi(\rightarrow \mu^+\mu^-)K^\pm$ is chosen for the measurement because the invariant mass of the decay products peaks narrowly at the mass of the B meson, the topology is similar than that of the signal process, and the branching fractions of the decays involved in the process are precisely known [3]. The events from the control process are reconstructed in data as well as in MC, from a dedicated sample simulating $B^\pm \rightarrow J/\psi(\rightarrow \mu^+\mu^-)K^\pm$ decays. Events are selected if the leading muon, fulfilling the *soft* identification criteria, has $p_T > 7$ GeV and $|\eta| < 1.5$, the subleading muon has $p_T > 2$ GeV and $|\eta| < 1.8$, and the kaon has $p_T > 1$ GeV and $|\eta| < 1.6$. The dimuon system and the kaon are fitted to a common vertex with $L_{xy} > 350 \mu\text{m}$, and the $K^\pm\mu^+\mu^-$ system must have $p_T > 3$ GeV.

The cross section of interest can be obtained as

$$\sigma_{B^\pm}^{\text{eff}} = N_{\text{control}}^{\text{data}} \times \frac{1}{\mathcal{L}} \times \frac{1}{\mathcal{B}(B^\pm \rightarrow J/\psi K^\pm)\mathcal{B}(J/\psi \rightarrow \mu^+\mu^-)} \times \frac{1}{\epsilon_{\text{control}}^{\text{MC}}}, \quad (6.20)$$

where

- $N_{\text{control}}^{\text{data}}$ is the number of control events in data, obtained from a fit to the $K^\pm\mu^+\mu^-$ invariant mass, as described below.
- $\mathcal{L}$ is the luminosity of the subset of the B-parking sample used for this study.
- $\mathcal{B}(B^\pm \rightarrow J/\psi K^\pm)$ and $\mathcal{B}(J/\psi \rightarrow \mu^+\mu^-)$ are the relevant branching fractions for the control process [3].
- $\epsilon_{\text{control}}^{\text{MC}}$ is the efficiency obtained from simulation. It accounts for the acceptance, reconstruction, selection, and trigger effects. To be able to extrapolate the measurement of $\sigma_{B^\pm}^{\text{eff}}$ using the control process to the signal process, it is important that the simulated control sample is produced using the same generator conditions as for the signal samples.

The number of control events, $N_{\text{control}}^{\text{data}}$, is obtained from data using an extended unbinned maximum likelihood fit to the $K^\pm\mu^+\mu^-$ invariant mass, shown in Fig. 6.17. Three components must be taken into account:

1. The signal peak, fitted with the convolution of a Gaussian function with a Breit-Wigner function.
2. The partially reconstructed background, arising from misreconstructed B meson decays, such as $B \rightarrow J/\psi + \text{hadrons}$, and fitted with an exponential function.
3. The combinatorial background, originating from spurious combinations of a J/ψ meson with an uncorrelated track, and fitted with a Crystal Ball function.

The contribution from the $B^\pm \rightarrow J/\psi\pi^\pm$ decay is not taken into account, as it was found to be negligible.




Figure 6.17: Distribution of the $K^\pm\mu^\pm\mu^\mp$ invariant mass. The signal peak is observed at the mass of the B_u meson. The orange, green, and red curves show the contributions from signal, composite background, and combinatorial background, respectively, while the blue curve shows the fit to signal plus background.

The results are obtained using a subset of the B-parking sample, corresponding to an integrated luminosity of 0.77 fb^{-1} . This measurement is dominated by the systematic uncertainties, discussed in Section 6.8, and not by the statistical uncertainty. The measured value is $\sigma_{B^\pm}^{\text{eff}} = (572.0 \pm 4.9 \text{ (stat)} \pm 85.8 \text{ (syst)}) \text{ } \mu\text{b}$.

A similar measurement was performed at CMS in a different fiducial phase space [131]. The measurement of $\sigma_{B^\pm}^{\text{eff}}$ was validated by repeating the measurement in the same fiducial phase space as in Ref. [131]. The results were found to agree within the uncertainty.

6.6.6 Signal shape parametrisation

In the analysis, the presence of heavy neutrinos in the data is searched for in the form of a peaking structure over the background in the $\ell^\pm\pi^\mp$ invariant mass distributions. In the following, we describe how the signal shape depends on the m_N and $c\tau_N$ hypotheses and how it is parametrised for the final extraction of the results.

An example of a fit to the $m(\ell^\pm\pi^\mp)$ distribution is shown in Fig. 6.18 for a signal with $m_N = 1 \text{ GeV}$ and $c\tau_N = 1000 \text{ mm}$, in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category of the dimuon channel. The shape is well modelled by a double-sided Crystal Ball function [47, 48]. This function has a Gaussian core, described by the mean, μ , and the resolution, σ , as well as two powerlaw tails, each of them being described by the parameter α , which modulates the distance from μ for which the powerlaw tail starts to be relevant, and the parameter n , which regulates the steepness of the tail.

Dedicated studies showed that the four parameters of the tails can be kept fixed for all signal hypotheses, and in all the analysis categories, listed in Table 6.5. In the dimuon channel, the




Figure 6.18: Distribution of $m(\ell^\pm \pi^\mp)$ for simulated events with $m_N = 1$ GeV and $c\tau_N = 1000$ mm in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel. The blue curve corresponds to the fit to the data using a double-sided Crystal Ball function.

tails are symmetric, while they are not in the mixed-flavour channel, stemming from the effect of Bremsstrahlung of electrons. Moreover, the resolution, σ , can be parametrised only in the mass, m_N , and be kept constant for the different $c\tau_N$ hypotheses and in the different categories. The residual effect is accounted for with a dedicated systematic uncertainty, as described in Section 6.8. The resolution, averaged over the different $c\tau_N$ hypotheses, is shown as a function of m_N in Fig. 6.19 in the inclusive category of the dimuon channel. The resolution is well described by a linear function as a function of m_N .

In the final fit, the only parameter that is not fixed is the mean, μ , which is left floating within the range corresponding to the lepton energy corrections, as discussed in Section 6.8.

6.7 Background estimation

The analysis is designed in such a way that the long-lived heavy neutrino decays exclusively into a charged lepton and a charged pion, thus allowing the reconstruction of its total invariant mass. The main sources of background passing the selection criteria are described in Section 6.7.1. In particular, several contributions from SM resonances must be vetoed, as discussed in Section 6.7.2. The strategy of the analysis then consists in a bump hunt in the invariant mass spectrum of the displaced objects. A direct implication of this strategy is that the background is estimated from a direct fit to the data, following the procedure explained in Section 6.7.3.

6.7.1 Background sources

The main signatures of the final states are three objects with low values of p_T and the presence of a displaced vertex. As discussed in Section 6.4, the vast majority of background events with such signatures are rejected by the preselection criteria and the pNN-based selection. Background




Figure 6.19: Resolution as a function of the signal mass m_N for events in the dimuon channel. For each mass point, the resolution is averaged over the given $c\tau_N$ hypotheses. The points are fitted with a polynomial function of first order.

events that nevertheless pass the selection requirements originate from low- p_T QCD processes that can be grouped in three main families of processes:

1. B cascade decays: genuine B meson decays that mimic the signal signatures, e.g. $B \rightarrow D$ decays.
2. Combinatorial background: events originating from the combination of uncorrelated tracks forming a signal candidate.
3. Instrumental background: events originating from pileup, secondary interactions in the tracker material, and other detector reconstruction effects.

The contribution of these three families of processes is shown in Fig. 6.20 for the $\mu^\pm\pi^\mp$ invariant mass and $L_{xy}/\sigma_{L_{xy}}$ spectra in the dimuon channel. The distributions are shown in the inclusive signal region for simulated events for muon-enriched low- p_T QCD processes that pass the preselection criteria, described in Section 6.4.1, except for the vetoes on the SM resonances. To assign the events to either of the three families of background processes, the reconstructed particles are matched to the corresponding generator-level particle, for which the chain of ancestors is saved. The matching is performed with the requirements that the two particles are within $\Delta R < 0.015$ of each other and that their transverse momenta agree within 10%. Such a matching is not possible for the background processes in the third family, as the associated generator-level particles are not saved in the samples. In Fig. 6.20 (left), we observe that the B cascade decays contribute up to a mass of about 3 GeV, which corresponds to the kinematic threshold of the $B \rightarrow D$ decay channel, after which the background arises uniquely from combinatorial and instrumental effects. In this spectrum, the presence of resonances at 1.77 GeV and 3.1 GeV is clearly visible, as further discussed below. In Fig. 6.20 (right), we verify that the




Figure 6.20: Distribution of the $\mu^\pm\pi^\mp$ invariant mass (*left*) and $L_{xy}/\sigma_{L_{xy}}$ (*right*) for simulated QCD background events in the dimuon channel. The contributions from B cascade decays, combinatorial background, and instrumental background are indicated with the magenta, green, and grey colour, respectively.

main sources of background at large $L_{xy}/\sigma_{L_{xy}}$ originate from combinatorial and instrumental processes.

6.7.2 Standard model resonances

A series of SM processes give rise to resonances in the different two-body invariant mass spectra. To avoid the contamination of these background processes in the signal regions, vetoes are applied on events for which the corresponding invariant masses are near the mass of a resonance. The list of resonances and the corresponding vetoes is given in Table 6.6 for the dimuon channel. The same vetoes are applied on the corresponding mass spectra in the mixed-flavour channel.

Three dimuon resonances are identified in the $\mu_B\mu^\pm$ invariant mass spectrum as coming from the decay of a $\phi(1020)$, J/ψ , and $\psi(2S)$. Because these particles are neutral, the associated vetoes are only applied for events in the OS categories. In the $\mu_B\pi^\mp$ invariant mass spectrum, events with a misidentified muon can contribute to the $D^0 \rightarrow K\pi$ decay, while events with a misidentified pion can contribute to the $J/\psi \rightarrow \mu\mu$ process. As the D^0 and J/ψ particles are neutral and the displaced muon, $\mu^\pm$, and pion, $\pi^\mp$, have opposite charges, the associated vetoes are applied only for events in the SS categories. For the D^0 decay, the vetoed mass window is not centred around the mass of the D^0 particle to account for the fact that the correct particle mass was not assigned to the final state tracks. Finally, three peaks were identified in the $\ell^\pm\pi^\mp$ invariant mass spectrum, two of them being clearly visible in Fig. 6.20 (*left*). They originate from the two-body decays of the D^0 , J/ψ , and $\psi(2S)$ particles and involve one misidentified particle in the final state. Because these resonances appear in the mass spectrum in which the presence of signal is searched for, no signal extraction nor exclusion limit setting are performed in the corresponding vetoed mass regions, such that the need for disentangling the contribution from an HNL signal and from the SM decay is avoided in case a peaking structure is found in

Table 6.6: List of SM resonances and corresponding vetoes in the $\mu_B\mu^\pm$ and $\mu_B\pi^\mp$ invariant mass spectra (*upper part*) and in the $\mu^\pm\pi^\mp$ invariant mass spectrum (*lower part*). For each process, the categories in which the veto is applied is reported and the presence of misidentified particles is indicated.

Mass spectrum	Process	Veto (GeV)	Categories	Misidentification
$m(\mu_B\mu^\pm)$	$\phi(1020) \rightarrow \mu\mu$	$ m(\mu_B\mu^\pm) - 1.02 > 0.01$	OS	0
	$J/\psi \rightarrow \mu\mu$	$ m(\mu_B\mu^\pm) - 3.10 > 0.15$	OS	0
	$\psi(2S) \rightarrow \mu\mu$	$ m(\mu_B\mu^\pm) - 3.69 > 0.08$	OS	0
$m(\mu_B\pi^\mp)$	$D^0 \rightarrow K\pi$	$ m(\mu_B\pi^\mp) - 1.76 > 0.05$	SS	1 misid. μ
	$J/\psi \rightarrow \mu\mu$	$ m(\mu_B\pi^\mp) - 3.10 > 0.05$	SS	1 misid. π
$m(\mu^\pm\pi^\mp)$	$D^0 \rightarrow K\pi$	$ m(\mu^\pm\pi^\mp) - 1.77 > 0.03$	all	1 misid. μ
	$J/\psi \rightarrow \mu\mu$	$ m(\mu^\pm\pi^\mp) - 3.10 > 0.05$	all	1 misid. π
	$\psi(2S) \rightarrow \mu\mu$	$ m(\mu^\pm\pi^\mp) - 3.70 > 0.05$	all	1 misid. π

data. The three vetoed regions listed in the lower part of Table 6.6 will therefore appear as blank regions in the limit plots presented in Section 6.9.

6.7.3 Background shape modelling

As further explained in Section 6.9, the signal of mass m_N is not searched for using the entire $m(\ell^\pm\pi^\mp)$ spectrum between 1 and 6 GeV, but rather in the mass window of size $\pm 10\sigma$ centred around m_N , where σ is the resolution introduced in Section 6.6.6. Prior to the final fit to extract the signal, the background shape is modelled from a direct fit to the data in the sidebands of the $m(\ell^\pm\pi^\mp)$ spectrum, defined as the regions outside the signal window of size $\pm 3\sigma$ around the signal mass point m_N . However, the functional form describing the background is not known a priori and may, in fact, be different across the mass windows and the event categories. To account for the uncertainty related to the choice of the functional form used to fit the data, the *discrete profiling* method is employed [132].

The method consists in considering not a single, but a set of functions, $\mathcal{S}$, providing a good description of the background shape, and in treating the choice of the function as a systematic uncertainty handled as a discrete nuisance parameter. During the final fit, the profile-likelihood is evaluated for each function and the best fit is chosen.

Three families of functions are considered for the fit to the data: (i) power-law polynomials, (ii) sum of exponentials, and (iii) Laurent-series polynomials. In principle, each family comprises all the functions of any order N . However, functions with high values of N may overfit the data. To avoid such an effect, merely functions with order $N_{\max}$ are considered as candidates to enter the final set of functions, $\mathcal{S}$. For each search category and mass window, the maximum order, $N_{\max}$, for each family is determined using a Fisher test [133]. This statistical procedure is useful to assess if the function of order $N+1$ provides a better description of the data than the function




Figure 6.21: Distribution of $m(\ell^\pm\pi^\mp)$ in the mass window around 2 GeV in the low $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. (Left) The coloured dashed lines represent the different functions of the set $\mathcal{S}$ obtained in the sidebands. The grey hashed region represents the signal region of size $\pm 3\sigma$. (Right) The red curve shows the result of the background-only fit to the data and the green curve shows an example of the expected Majorana signal with $m_N = 2$ GeV and $c\tau_N = 10$ mm for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$.

of order N . In practice, the successive orders of a function are tested iteratively according to three steps:

1. If N is the lowest order, or N was found in the previous iteration to describe well the data, the function of order $N + 1$ is tested by computing $\Delta\text{NLL} = 2(\text{NLL}_N - \text{NLL}_{N+1})$, where NLL_N is the negative log-likelihood of the fit with the function of order N , augmented by a penalty term equal to half of the number of the function parameters. The penalty term, which scales with larger values of N , is applied as a protection against the risk of overfitting the data.
2. The p -value, p^F , of the fit is computed as $p^F = p(\chi^2 > \Delta\text{NLL})$, which is possible as ΔNLL behaves asymptotically as a χ^2 distribution, according to Wilk's theorem [134].
3. If $p^F < 0.05$, the function of order $N + 1$ is supported by the data, and the test is carried out for the next-order function. Otherwise, the next-order function is not tested and $N_{\text{max}} = N + 1$.

After N_{max} is determined, the function with order N_{max} is added to the set of functions, $\mathcal{S}$, as well as the functions with order $N < N_{\text{max}}$, provided that their associated χ^2 -goodness of fit probability is greater than 0.01.

An example of the outcome of the background shape modelling procedure is shown in Fig. 6.21 (left) for data in the mass window of size $\pm 10\sigma$ around 2 GeV in the low $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. In this case, $N_{\text{max}} = 1$ for the three functional

Table 6.7: List of systematic uncertainties impacting the expected signal event yield. The ranges in the value represent the uncertainties across the different event categories.

Source	Value (%)
Signal shape	15
$\sigma_{B^\pm}^{\text{eff}}$	15
f_c	24
Signal selection	5–20
Matching	5
Tracking efficiency	5
Trigger scale factors	5
Muon identification scale factors	1
Electron identification scale factors	3

forms. We verify that all functions can fit the data well, while being distinct from one another. The corresponding background-only fit to the data is shown in Fig. 6.21 (*right*), in which an example of a Majorana signal with $m_N = 2$ GeV and $c\tau_N = 10$ mm is overlaid for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$.

6.8 Systematic uncertainties

Several sources of systematic uncertainties must be taken into account when fitting the data to extract the signal and to set exclusion limits. Generally speaking, these uncertainties are small compared to the statistical uncertainties arising from the limited size of the event sample.

As discussed in Section 6.7.3, the uncertainty in the signal event yield associated with the modelling of the background shape is accounted for using the discrete profiling method. In the final fit, the associated parameters are free-floating and are treated as uncorrelated in the different event categories and the different mass windows.

The remaining sources of uncertainties are related to the parametrisation of the signal shape and to the normalisation of the signal event yield. They are listed in Table 6.7 and are discussed in the following.

Signal shape parametrisation

As discussed in Section 6.6.6, the signal shape is parametrised with a double-sided Crystal Ball function, using fixed parameters across the 24 event categories and for the different $c\tau_N$ hypotheses. A systematic uncertainty of 15% is applied to account for the possible variation of the signal shape across the event categories and the range of $c\tau_N$ values. This uncertainty was obtained from the relative difference in the expected median limit for representative signal points when varying the parameters of the functional form.

Additionally, as anticipated in Section 6.6.6, the mean of the signal distribution is set to the probed mass m_N in the final fit, with a freedom of 0.1% and 2% for signals in the dimuon and mixed-flavour channels, respectively, to account for the lepton energy uncertainties. The other

parameters, i.e. the resolution and the parameters of the tails, are kept fixed.

Signal normalisation

The other systematic uncertainties are associated with the number of expected signal events, obtained using Eq. (6.4). First of all, we note that the luminosity, $\mathcal{L}$, appears in the numerator in this equation, while it appears in the denominator in Eq. (6.20) for the measurement of $\sigma_{B^\pm}^{\text{eff}}$. Therefore, the associated uncertainty mostly cancels out and does not impact the final results. On the other hand, the uncertainties affecting the respective efficiencies, ϵ_{sig}^q and $\epsilon_{\text{control}}^{\text{MC}}$, do not cancel as the measurement of $\sigma_{B^\pm}^{\text{eff}}$ was not performed using the full luminosity of the B-parking data set.

The uncertainty in the measurement of the effective $B^\pm$ cross section, $\sigma_{B^\pm}^{\text{eff}}$, is estimated to be 15%. The main contributions to this uncertainty are related to the muon trigger, to the model used in the likelihood fits, and to the reconstruction of the B decay vertex.

Moreover, for events in the high $m(\ell_B \ell^\pm \pi^\mp)$ categories, to which only signals originating from the decay of B_c mesons contribute, an uncertainty of 24% is applied in the fragmentation fraction f_c , as obtained from the measurement presented in Ref. [121].

Finally, the uncertainties related to the signal efficiency, ϵ_{sig}^q , in Eq. (6.4), are

1. The uncertainty arising from the pNN selection, described in Section 6.4.3. This uncertainty was obtained using the well-defined SM process, $B^\pm \rightarrow J/\psi(\rightarrow \mu^+ \mu^-) K^\pm$. After the events are reconstructed and selected, the pNN is evaluated on both data and MC events with a mass key providing a signal-like behaviour. An example of the resulting pNN score distributions in data and MC is shown in Fig. 6.22 for the high $L_{xy}/\sigma_{L_{xy}}$, OS and low $m(\ell_B \ell^\pm \pi^\mp)$ categories in the dimuon channel. The ratio of the efficiency in data and MC of the selection on the pNN score to be greater than 0.99, which is the threshold used to select signal events for the final fit, is used as an uncertainty. Depending on the event category, it ranges from 5 to 20%.
2. The uncertainty of 5% associated with the matching, discussed in Section 6.3, between a reconstructed candidate to the underlying generated candidate.
3. The uncertainty covering the effect of potential mismodelling of quantities related to the tracking or vertexing associated with displaced decay vertices. This uncertainty was studied in the context of Ref. [107] and was found to be at most 5%. This value is conservatively applied across all the event categories.
4. The uncertainty arising from the measurement of the corrections for the trigger efficiency, discussed in Section 6.5. An uncertainty of at most 5% was obtained by studying the impact on the expected number of signal events for representative signal hypotheses when varying the corrections within one standard deviation. A systematic uncertainty of 5% is conservatively applied across all the event categories and for all signal hypotheses.




Figure 6.22: Distribution of score of the pNN evaluated in the $B^\pm \rightarrow J/\psi(\rightarrow \mu^+\mu^-)K^\pm$ channel for data and MC events in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B \ell^\pm \pi^\mp)$ category in the dimuon channel. The pNN is evaluated with a mass key equal to 4.5 GeV.

5. The uncertainties in the efficiency of the lepton identification criteria, discussed in Section 6.5. The uncertainty for the *soft* and *loose* muon identification is estimated to be below 1% using a similar method as for the uncertainty in the corrections for trigger efficiency discussed above. The measurement of the uncertainty in the electron identification efficiency was not part of the thesis work.

The systematic uncertainties are modelled as log-normal distributions. Moreover, they are treated as uncorrelated across the different analysis categories in the final fit, except for the uncertainties in $\sigma_{B^\pm}^{\text{eff}}$ and f_c .

6.9 Results and interpretation

As the value of the heavy neutrino mass, m_N , is not predicted, the presence of a new physics signal is probed in data with a scan on masses in the range $1 < m_N < 6$ GeV with unprecedented resolution. Indeed, the step size between the trial masses is approximately twice the resolution, σ , defined in Section 6.6.6, implying that more than a hundred mass hypotheses are tested in the given mass range, provided that they are not within one of the vetoed regions listed in Table 6.6. For a given mass m_N , the events are selected according to the preselection criteria described in Section 6.4.1 and with the requirement on the score of the pNN, evaluated for that given mass m_N , to be greater than 0.99, as discussed in Section 6.4.3. The signal event yield is then extracted from a parametric fit to the $m(\ell^\pm \pi^\mp)$ distributions in data in a window of size $\pm 10\sigma$ around the tested mass m_N . The fit is performed simultaneously in the 24 event categories listed in Table 6.5 and uses the sum of signal plus background models, discussed in

Sections 6.6.6 and 6.7.3 for the signal and background components, respectively. Moreover, the choice of the background functional form is treated as a nuisance parameter, as well as the other sources of systematic uncertainties presented in Section 6.8.

No significant excesses over the SM background predictions were observed in any of the mass windows. Examples of fits to the $m(\ell^\pm\pi^\mp)$ distributions are shown in Fig. 6.23 for signal masses $m_N = 1.0, 1.5, 2.0$, and 2.5 GeV for events in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. The fits correspond to signal hypotheses with $c\tau_N = 10000, 2000, 700$, and 100 mm, respectively. These values are chosen as they are close to the exclusion limit in the case where the HNL mixes exclusively with muon neutrinos, as further discussed below. In these examples, the distributions of data in the mass windows centred at $m_N = 1.0, 1.5$, and 2.0 GeV do not present any signal-like structure, and are representative of the $\mathcal{O}(1000)$ mass spectra fitted in the search. Instead, the distribution of data in the mass window centred at $m_N = 2.5$ GeV shows a signal-like fluctuation. For the Majorana case, the fluctuation has a local significance of 2.6σ when combining the results from each of the OS and SS categories, while for the Dirac-like case, the local significance is 2.8σ when combining the results from each of the OS categories.

With the absence of a clear sign for the presence of heavy neutrinos in the data, the results of the search are used to set constraints on the parameter space describing the heavy neutrino phenomenology. In the following, exclusion limits on $|V_N|^2$ and $c\tau_N$ are presented at 95% confidence level (CL) as a function of the HNL mass, m_N , for multiple signal hypotheses, specified below. The statistical procedure to obtain the limits is summarised in Section 4.8. For each signal hypothesis, an unbinned maximum likelihood function, $\mathcal{L}(r, \vec{\theta})$, is constructed with the signal strength, r , as the parameter of interest and with the set of nuisance parameters, $\vec{\theta}$, mentioned above. The limits are computed in the asymptotic approximation [79] using the CL_s criterion [80] and are obtained with the combination of the different event categories. A given signal hypothesis is excluded in case the condition $r \geq 1$ is not compatible with the observed data at 95% CL.

The various signal hypotheses scrutinised in the search are specified by (i) the value of m_N ; (ii) the value of $c\tau_N$ or equivalently of $|V_N|^2$, depending on which quantity the limits are set; (iii) the mixing scenario defined by the value of the mixing ratios (r_e, r_μ, r_τ) ; and (iv) the nature of the heavy neutrino, which can be either Majorana or Dirac-like. For each scenario, the associated expected signal event yield is obtained using Eq. (6.4). The yield is reweighted to any mixing scenario using the method described in Section 6.6.2, while it is reweighted to the Dirac-like case following the procedure described in Section 6.6.3. Moreover, for each value of m_N , a fine scan on the value of $c\tau_N$, or equivalently $|V_N|^2$, is performed using the $c\tau_N$ -reweighting method described in Section 6.6.4.

Exclusion limits are computed for four mixing scenarios. The first one considers the case in which the heavy neutrino mixes exclusively with the muon sector, i.e. $(r_e, r_\mu, r_\tau) = (0, 1, 0)$. It is referred to as the muon-exclusive mixing scenario. Constraining this scenario is useful as it has been considered as a benchmark scenario for virtually all the previous HNL searches involving muons in the final state, thus allowing for a direct comparison of our results with




Figure 6.23: Fits to the $m(\ell^\pm\pi^\mp)$ distributions for signals with mass $m_N = 1.0$ GeV (upper left), 1.5 GeV (upper right), 2.0 GeV (lower left), and 2.5 GeV (lower right), in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category of the dimuon channel. The signal hypotheses correspond to $c\tau_N = 10000$, 2000, 700, and 100 mm, respectively. The blue curve shows the signal-plus-background fit, while the green and red curves show the individual signal and background components, respectively. The yellow band corresponds to the total systematic plus statistical uncertainty.

those from the previous searches. The three other benchmark scenarios, proposed in Ref. [117], are $(r_e, r_\mu, r_\tau) = (0, 1/2, 1/2)$, $(1/2, 1/2, 0)$, and $(1/3, 1/3, 1/3)$. They are referred to as the mixed-flavour mixing scenarios given that the value of at least two ratios is non-null. The limits for the first scenario are obtained using the categories in the dimuon channel only, while those for the other two scenarios are obtained using the categories in both the dimuon and mixed-flavour channels. The results for the latter channel were derived by the members of the INFN and Sapienza institutes, while the task of combining the results from the two channels was an integral part of the thesis work. Finally, the results are interpreted as limits on $|V_N|^2$ and $c\tau_N$ for 66 different mixing scenarios, in the form of ternary plots. The interpretation of the results for each mixing scenario is discussed below.




Figure 6.24: Expected and observed 95% CL limits on $|V_N|^2$ as a function of m_N in the Majorana (left) and Dirac-like (right) scenarios. The limits are shown for the muon-exclusive mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$. Results from the CMS [107, 113], ATLAS [105], LHCb [103], and Belle [102] Collaborations are reported. The grey bands indicate the mass regions that are vetoed, as listed in Table 6.6.

Muon-exclusive mixing scenario

The observed and expected upper exclusion limits at 95% CL on $|V_N|^2$ are shown in Fig. 6.24 as a function of m_N for the Majorana and Dirac-like scenarios. The limits for $|V_N|^2 > 10^{-2}$ are not shown, as such strongly-coupling heavy neutrinos are discarded by well-established SM measurements [135]. The best limit is obtained for an HNL mass $m_N = 1.95$ GeV for the Majorana scenario and for an HNL mass $m_N = 1.68$ GeV for the Dirac-like scenario, for which values of $|V_N|^2$ are excluded down to 2.0×10^{-5} and 3.2×10^{-5} , respectively.

The limits for lower signal masses slightly degrade, as a consequence of a loss in the CMS signal acceptance due to the constraint that the heavy neutrino must decay within the tracker volume. It can indeed be deduced from Eq. (5.18) that, for a fixed value of $|V_N|^2$, signals with smaller masses have a longer lifetime, and are therefore more likely to decay beyond the tracker. A similar argument explains the loss in the search sensitivity at higher signal masses. For a fixed value of $|V_N|^2$, the lifetime of signals with larger masses are shorter. Consequently, as the present search is optimised for long-lived signatures, its performance is affected for short-lived signal hypotheses. Furthermore, for these large masses, the loss in the sensitivity also follows from the decreasing values of the HNL production and decay branching fractions, as seen in Fig. 6.4 (right) and Fig. 6.6.

Moreover, structures in the limits are present for masses approximately in the range $3 < m_N < 4$ GeV. The main cause driving this effect comes from the transition between regimes in

which the B meson decays semileptonically or leptonically, as shown in Fig. 6.4 (left).

For the Majorana scenario, the SS event categories generally drive the sensitivity of the search, as the background event yield tends to be smaller in these categories than in the OS categories, because of the violation of the lepton number. Moreover, the high, medium, and low $L_{xy}/\sigma_{L_{xy}}$ categories contribute the most to the signal exclusion for signal masses between approximately 1 and 3 GeV, 3 and 5 GeV, and 5 and 6 GeV, respectively. Finally, for the latter mass range, the signal arises uniquely from the decays of B_c mesons, and the low $m(\ell_B \ell^\pm \pi^\mp)$ categories hence do not contribute. For masses in the range $3 < m_N < 5$ GeV, the addition of the high $m(\ell_B \ell^\pm \pi^\mp)$ categories improves the limits by up to about 40%.

The sensitivity for the Dirac-like scenario is smaller than that in the Majorana scenario. Although the expected signal yield, for a given signal hypothesis, in the OS categories is twice as large, as discussed in Section 6.6.3, in the former scenario than in the latter, it is not sufficient, in the context of this search, to outweigh the loss in sensitivity resulting from the null yields in the SS categories.

Limits obtained from results of previous HNL searches probing masses in the same range are overlaid in Fig. 6.24. These are

1. Limits obtained by the ATLAS [105] and CMS [107, 113] Collaborations. These searches probe long-lived HNLs in the decays of W bosons and exploit about 138 fb^{-1} of data, i.e. more than three times as much data as for the present search. The limits in the present analysis are improved by up to a factor of about three for $m_N < 1.75$ GeV for the Majorana scenario, and a factor of about two for $m_N < 1.7$ GeV for the Dirac-like scenario.
2. Limits for the Majorana scenario obtained by the Belle [102] and LHCb [103] (and revised in Ref. [136]) Collaborations. These searches target a similar signal process as the present search, namely long-lived HNLs in the decays of B mesons, using 711 fb^{-1} and 3 fb^{-1} of data, respectively. Compared to the Belle analysis, the limits are improved by up to a factor of about two, while compared to the LHCb analysis, the limits are improved by up to one order of magnitude for all mass hypotheses.

As a consequence, the upper exclusion limits on $|V_N|^2$ for $1 < m_N < 1.7$ GeV are the most stringent to date at a collider experiment. We note that beam-dump experiments, such as the CHARM [100] and BEBC [137, 138] experiments, were sensitive to signal masses up to about 2 GeV. However, these searches do not present upper bounds on the excluded value of $|V_N|^2$ in a way that allows a straightforward comparison with the sensitivity of this analysis.

Mixed-flavour mixing scenarios

Figure 6.25 shows the observed and expected upper exclusion limits at 95% CL on $|V_N|^2$ obtained for the three mixed-flavour scenarios, $(r_e, r_\mu, r_\tau) = (0, 1/2, 1/2)$, $(1/2, 1/2, 0)$, and $(1/3, 1/3, 1/3)$ for both the Majorana and Dirac-like scenarios. The value of the best exclusion limits on $|V_N|^2$ for each scenario is summarised in Table 6.8. The sensitivity for these scenarios is less smaller than for the muon-exclusive mixing scenario, as can be inferred from Eq. (6.13). The




Figure 6.25: Expected and observed 95% CL limits on $|V_N|^2$ as a function of m_N in the Majorana (left) and Dirac-like (right) scenarios. The limits are shown for the mixed-flavour mixing scenarios $(r_e, r_\mu, r_\tau) = (0, 1/2, 1/2)$ (upper), $(r_e, r_\mu, r_\tau) = (1/2, 1/2, 0)$ (middle), and $(r_e, r_\mu, r_\tau) = (1/3, 1/3, 1/3)$ (lower). Results from the CMS [112] Collaboration are reported. The grey bands indicate the mass regions that are vetoed, as listed in Table 6.6.

Table 6.8: Most stringent observed upper limits on $|V_N|^2$ at 95% CL for the different mixing scenarios. The minimum excluded value of $|V_N|^2$ is reported together with the signal mass at which it is obtained.

(r_e, r_μ, r_τ)	Scenario	$ V_N ^2$	Mass (GeV)
(0, 1, 0)	Majorana	2.0×10^{-5}	1.95
(0, 1/2, 1/2)	Majorana	4.0×10^{-5}	1.42
(1/2, 1/2, 0)	Majorana	3.3×10^{-5}	2.15
(1/3, 1/3, 1/3)	Majorana	5.0×10^{-5}	2.15
(0, 1, 0)	Dirac-like	3.2×10^{-5}	1.68
(0, 1/2, 1/2)	Dirac-like	6.5×10^{-5}	1.68
(1/2, 1/2, 0)	Dirac-like	5.7×10^{-5}	1.68
(1/3, 1/3, 1/3)	Dirac-like	8.5×10^{-5}	1.68

smaller sensitivity does not only manifest itself in the larger excluded values of $|V_N|^2$, but also in the fact that the limit closes for masses around 4 GeV. For larger masses, the heavy neutrino is in fact produced only via leptonic decays of B mesons, as seen in Fig. 6.4 (left), and the associated expected signal yield is thus smaller compared to that obtained for smaller masses at which the semileptonic decays of B mesons are the major contribution.

A previous search from CMS [112] presented limits for the mixed-flavour mixing scenarios for $m_N > 2$ GeV. This analysis provides the first limits for these scenarios for masses in the range $1 < m_N < 2$ GeV.

Ternary plots

Finally, a scan on different mixing scenarios (r_e, r_μ, r_τ) , with a step size of 0.1, was performed for HNL masses $m_N = 1.0, 1.5$, and 2.0 GeV. The constraint $r_e + r_\mu + r_\tau = 1$ reduces the number of possible scenarios to 66 and allows the results to be displayed in the form of ternary plots. For each mixing scenario, the value of the upper (lower) observed exclusion limit at 95% CL on $|V_N|^2$ ($c\tau_N$) is reported. The results are obtained with the combination of the dimuon and mixed-flavour categories. It is the first time that such plots are presented for $m_N \leq 2$ GeV.

Figure 6.26 shows the observed upper limit on $|V_N|^2$ for the 66 different mixing scenarios, for the three HNL masses, and for both the Majorana and Dirac-like cases. The red crosses indicate that no limit is obtained for the scenarios with $r_\mu = 0$, given that the final state of the search must always contain at least one muon. We verify that the value of the limit on $|V_N|^2$ gets smaller with increasing values of r_μ , reaching the best performance for $r_\mu = 1$. There is also such a trend with increasing values of r_e , although it is less pronounced than for the muon mixing ratio. The values of the best limits on $|V_N|^2$, obtained for the muon-exclusive mixing scenario, are reported in Table 6.9 for the three signal masses, and the Majorana and Dirac-like scenarios.

The limits on $|V_N|^2$ are related to those on $c\tau_N$ according to the relation in Eq. (6.10). The ternary plots showing the observed lower exclusion limits on $c\tau_N$ are presented in Fig. 6.27 for the three mass hypotheses, and for the Majorana and Dirac-like cases. The best limits on $c\tau_N$ are




Figure 6.26: Observed 95% CL upper limits on $|V_N|^2$ as functions of the mixing ratios (r_e , r_μ , r_τ) for fixed HNL masses of 1 GeV (*upper*), 1.5 GeV (*middle*), and 2 GeV (*lower*), in the Majorana (*left*) and Dirac-like (*right*) scenarios. The range of values on the z -axis is not kept fixed across the different plots. The red crosses indicate that there is no exclusion for that point. The orientation of the value markers on each axis indicates the direction of reading.




Figure 6.27: Observed 95% CL lower limits on $c\tau_N$ as functions of the mixing ratios (r_e , r_μ , r_τ) for fixed HNL masses of 1 GeV (upper), 1.5 GeV (middle), and 2 GeV (lower), in the Majorana (left) and Dirac-like (right) scenarios. The range of values on the z -axis is kept fixed across the different plots. The red crosses indicate that there is no exclusion for that point. The orientation of the value markers on each axis indicates the direction of reading.

Table 6.9: Summary of the most stringent upper limits at 95% CL on $|V_N|^2$, obtained for the muon-exclusive mixing scenario. The minimum excluded value of $|V_N|^2$ is reported for masses $m_N = 1.0, 1.5$ and 2.0 , and for the Majorana and Dirac-like scenarios.

Mass (GeV)	Scenario	$ V_N ^2$
1.0	Majorana	5.9×10^{-5}
1.5	Majorana	3.1×10^{-5}
2.0	Majorana	2.6×10^{-5}
1.0	Dirac-like	1.0×10^{-4}
1.5	Dirac-like	6.5×10^{-5}
2.0	Dirac-like	4.3×10^{-5}

Table 6.10: Summary of the most stringent lower observed limits at 95% CL on $c\tau_N$, obtained for the muon-exclusive mixing scenario. The maximum excluded value of $c\tau_N$ is reported for masses $m_N = 1.0, 1.5$ and 2.0 , and for the Majorana and Dirac-like scenarios.

Mass (GeV)	Scenario	$c\tau_N$ (m)
1.0	Majorana	9.2
1.5	Majorana	2.3
2.0	Majorana	0.7
1.0	Dirac-like	10.5
1.5	Dirac-like	2.2
2.0	Dirac-like	0.8

obtained for $m_N = 1$ GeV in the muon-exclusive mixing scenario and correspond to $c\tau_N < 9.2$ m and $c\tau_N < 10.5$ m for the Majorana and Dirac-like cases, respectively. The most stringent limits on $c\tau_N$ obtained for all masses are summarised in Table 6.10.

While the limits improve with increasing values of r_μ , they slightly degrade with increasing values of r_e , as opposed to the limits on $|V_N|^2$. This behaviour is explained by the significantly smaller values, as shown in Fig. 6.6, of $\tilde{\Gamma}_\tau$ compared to $\tilde{\Gamma}_e$ and $\tilde{\Gamma}_\mu$ in Eq. (6.10). Moreover, unlike the limits on $|V_N|^2$ discussed above, the limits on $c\tau_N$ are more stringent for the Dirac-like case than for the Majorana one. This trend is a consequence of the reduced widths, $\tilde{\Gamma}$, in Eq. (6.10) being twice as low for the Dirac-like scenario than for the Majorana one.

6.10 Summary

A search for long-lived heavy neutrinos was conducted, for the first time at CMS, in the leptonic and semileptonic decays of B mesons produced in proton-proton collisions at $\sqrt{s} = 13$ TeV. The search is made possible by the collection of the B-parking data set, corresponding to an integrated luminosity of 41.6 fb^{-1} and accumulated using single-muon triggers designed to capture the signatures of the decays of beauty mesons. The search probes final states with two leptons, which can either be a muon or an electron, and a pion; at least one lepton must be a muon that satisfies the trigger requirements. As the HNL is assumed to be long lived, one lepton must originate, together with the pion, from a vertex displaced with respect to the luminous region. A search is performed for a peak in the invariant mass spectrum of the two displaced particles,

allowing the probe over a fine grid of HNL masses. The sensitivity of the search to the different mass hypotheses is optimised using a multivariate algorithm. No significant sign for the presence of new physics in the data is found. Exclusion limits at 95% CL on $|V_N|^2$ and $c\tau_N$ are derived in the mass range $1 < m_N < 6$ GeV for the scenario in which the heavy neutrino mixes exclusively with the muon neutrinos, as well as for scenarios in which it mixes with more than one neutrino flavour family, using either the Majorana or Dirac-like hypothesis. The most stringent exclusion upper limits on $|V_N|^2$ for masses $1 < m_N < 1.7$ GeV from a collider experiment to date were obtained. Moreover, the first upper limits on $|V_N|^2$ for masses $1 < m_N < 2$ GeV are presented for the mixing scenarios $(r_e, r_\mu, r_\tau) = (0, 1/2, 1/2)$, $(1/2, 1/2, 0)$, and $(1/3, 1/3, 1/3)$. Finally, lower limits on $c\tau_N$ in the form of ternary plots are presented for the first time for masses $m_N \leq 2.0$ GeV. This work is published in the *Journal of High Energy Physics* [119] and is part of two review papers published in the *Physics Reports* journal [34, 120].

Conclusion

The eagerness to find answers to unsolved questions in particle physics motivated the research work presented in this thesis. Signs for new physics were searched for as the manifestation of new particles in the data collected by the CMS experiment at the CERN LHC.

New physics was first probed in the form of supersymmetric particles. Such hypothetical particles have been gathering a lot of interest over the past decades, notably as some of them are candidates for dark matter, and as the framework in which they are predicted addresses theoretical limitations of the SM, such as the hierarchy problem. In the thesis work, the presence of these particles was probed in final states with jets and large missing transverse momentum, leveraging on the high signal-to-background discrimination power of the M_{T2} variable. Particular highlights of the work were the study of the two main sources of background and the combination of all the Run 2 data sets to extract the results. No significant excess of the number of observed data events over the background prediction was observed. Exclusion limits on the parameters of several simplified SUSY models were derived and are the most stringent to date.

New physics was then searched for in the neutrino sector. The prediction of the existence of heavy neutrinos can simultaneously solve the puzzle on the mass of the SM neutrinos, the asymmetry between matter and antimatter, and the mystery on the nature of dark matter. A novel branch of searches for heavy neutrinos at CMS was established using the decays of B mesons. It is the first search for new physics that exploits the B-parking data set, which is a sample enriched with $b\bar{b}$ events. Besides the entire design of the analysis in the novel production mode, a challenging aspect of the work was the assumption of the hypothetical particle to be long lived, giving rise to a decay vertex displaced with respect to the luminous region. While no sign for new physics was found in the data, the results show not only that such a search is possible at CMS, but also that it provides the most stringent constraints to date on given regions of the parameter space.

For searches for new physics to be performant, solid particle reconstruction algorithms are essential. Another project presented in the thesis consisted in the optimisation of the ECAL clustering algorithm in view of preserving stable performance during the Run 3 of data taking, characterised by an increased level of noise in the detector. The work resulted in optimal performance in terms of noise rejection, signal preservation, and energy resolution, and is beneficial to all the CMS analyses using Run 3 data and relying on photons, electrons, jets, or missing transverse momentum.

The work conducted during the PhD thesis resulted in significant contributions to the physics program at the CMS experiment and to the progress of research in particle physics. It constitutes yet another step forward in the endeavour of particle physicists to find a path towards excellence.

Appendix A

Performance of the clustering optimisation

The optimisation of the clustering algorithm of the electromagnetic calorimeter is presented in Section 2.5. In this appendix, additional figures assessing the performance of the tuning of the energy thresholds are provided.

Figures A.1, A.2, and A.3 show, in bins of energy and $|\eta|$, the relative performance of the tuned and reference energy thresholds, assessed in terms of the resolution, efficiency, and noise rate, respectively. The two sets of thresholds as well as the figures of merit are defined in Section 2.5. The performance is shown for simulated photons interacting with the electromagnetic calorimeter with conditions corresponding to a total integrated luminosity of 450 fb^{-1} .

We observe that the resolution is generally improved by the optimisation, particularly in the low-energy and high- $|\eta|$ regions. For instance, the resolution is improved by about 50% for photons with an energy between 60 and 80 GeV and a pseudorapidity $|\eta| > 2.8$. The efficiency is also well preserved in all the bins, except in bins with $|\eta| > 2.60$, in which the more stringent algorithm parameters discard signal photons. However, the loss of efficiency is largely counterbalanced by the decrease of the noise rate. Indeed, up to 100% of the clusters originating from noise are rejected with the tuned clustering algorithm.




Figure A.1: Relative difference of the resolution obtained with the tuned and reference energy thresholds, shown in different bins of the simulated photon energy and $|\eta|$. The performance is shown for simulated photons interacting with the ECAL with conditions corresponding to an integrated luminosity of 450 fb⁻¹. The red curve is a reference corresponding to a transverse energy of 5 GeV. The blank bins with a hash symbol are not populated by clusters as the energy is below the energy thresholds.




Figure A.2: Relative difference of the efficiency obtained with the tuned and reference energy thresholds, shown in different bins of the simulated photon energy and $|\eta|$. The performance is shown for simulated photons interacting with the ECAL with conditions corresponding to an integrated luminosity of 450 fb^{-1} . The red curve is a reference corresponding to a transverse energy of 5 GeV. The blank bins with a hash symbol are not populated by clusters as the energy is below the energy thresholds.




Figure A.3: Relative difference of the noise rate obtained with the tuned and reference energy thresholds, shown in different bins of the cluster energy and $|\eta|$. The performance is shown for simulated photons interacting with the ECAL with conditions corresponding to an integrated luminosity of 450 fb⁻¹. The red curve is a reference corresponding to a transverse energy of 5 GeV. The blank bins with a hash symbol are not populated by clusters as the energy is below the energy thresholds. The dash symbol indicates bins in which the noise rate is null with the reference thresholds.

Statement of originality

The research presented in the thesis is the result of personal work. In particular, all the figures were produced by the author, unless stated otherwise in the caption. The work was performed within the CMS Collaboration at CERN, and would thus not be possible without the previous generations of physicists who thought, built, and maintained the experiment and its software.

The clustering optimisation, presented in Section 2.5, was designed and conducted by the author, with useful advice from Maria Giulia Ratti and Mauro Donegà. The scripts used to simulate the events were co-developed with Maria Giulia Ratti, while the rest of the scripts were written by the author. These scripts were shared with the rest of the collaboration and an internal note, of which the author is the primary editor, was written to facilitate the repetition of the optimisation for future runs of data taking. The results from this work are published in Ref. [49], of which the author is the primary editor.

The M_{T2} search, described in Chapter 4, is the result of many years of development from numerous contributors, both at ETH Zürich and at U.S. institutes at UCSB and UCSD. The event samples were primarily produced by Bennett Marsh, Mario Masciovecchio, and Maria Giulia Ratti, while existing analysis scripts were streamlined and adapted by the author to obtain results using the full Run 2 data. The analysis is described in an internal CMS analysis note, of which the author is a secondary editor, and the results are published in Ref. [68]. The paper editing and the journal publication process were not part of the thesis work.

The search for heavy neutrinos in B meson decays, presented in Chapter 6, is a novel search fully designed and led by the author, with the guidance of Riccardo Manzoni. The simulation framework was co-developed with Maria Giulia Ratti, while the frameworks for the sample processing and the analysis were created by the author. The strategy and technical aspects of the analysis were established during the thesis work from beginning to end, and the supervision of the work in the mixed-flavour channel, taken care of by members of Italian institutes at Sapienza and INFN Rome, was also part of the author's responsibilities. The analysis is described in an internal CMS analysis note, of which the author is the main editor. The review and publication processes were conducted by the author, and resulted in three publications: the main analysis paper [119], of which the author is the primary editor, and two review papers on B-parking techniques [34], and on CMS searches for heavy neutrinos [120], to which the author contributed.

Finally, the author was given the responsibility to present several CMS results at conferences, resulting in two publications. The first one [114] summarises searches for long-lived particles, and the second one [118] provides an overview of results from new physics searches.

List of acronyms

AUC	<u>A</u> rea <u>u</u> nder the <u>c</u> urve
CERN	<u>C</u> onseil <u>e</u> uropéen pour la <u>r</u> echerche <u>n</u> ucléaire
CMS	<u>C</u> ompact <u>m</u> uon <u>s</u> olenoid
CL	<u>C</u> onfidence <u>l</u> evel
CR	<u>C</u> ontrol <u>r</u> egion
CSC	<u>C</u> athode <u>s</u> trip <u>c</u> hamber
DF	<u>D</u> ifferent- <u>f</u> lavour
DT	<u>D</u> rift <u>t</u> ube
DV	<u>D</u> isplaced <u>v</u> ertex
ECAL	<u>E</u> lectromagnetic <u>c</u> alorimeter
HCAL	<u>H</u> adron <u>c</u> alorimeter
HLT	<u>H</u> igh <u>l</u> evel <u>t</u> rigger
HNL	<u>H</u> eavy <u>n</u> eutral <u>l</u> epton
ISR	<u>I</u> nitial <u>s</u> tate <u>r</u> adiation
L1	<u>L</u> evel <u>1</u>
LEP	<u>L</u> arge <u>e</u> lectron <u>p</u> ositron
LHC	<u>L</u> arge <u>h</u> adron <u>c</u> ollider
LO	<u>L</u> eading <u>o</u> rders
LSP	<u>L</u> ightest <u>s</u> upersymmetric <u>p</u> article
MC	<u>M</u> onte <u>C</u> arlo
MSSM	<u>M</u> inimal <u>s</u> upersymmetric <u>s</u> tandard <u>m</u> odel
NLO	<u>N</u> ext to <u>l</u> eading <u>o</u> rders
NN	<u>N</u> eural <u>n</u> etwork
OS	<u>O</u> pposite <u>s</u> ign
PDF	<u>P</u> arton <u>d</u> istribution <u>f</u> unction
PF	<u>P</u> article- <u>f</u> low
pNN	<u>P</u> arametric <u>n</u> eural <u>n</u> etwork
PV	<u>P</u> rimary <u>v</u> ertex
QCD	<u>Q</u> uantum <u>c</u> hromod <u>ynamics</u>
RPC	<u>R</u> esistive <u>p</u> late <u>c</u> hamber
SF	<u>S</u> ame- <u>f</u> lavour
SM	<u>S</u> tandard <u>m</u> odel
SR	<u>S</u> ignal <u>r</u> egion
SS	<u>S</u> ame <u>s</u> ign
SUSY	<u>S</u> upers <u>ymmetry</u>
ν MSM	Neutrino (ν) <u>m</u> inimal <u>s</u> tandard <u>m</u> odel

List of figures

1.1	Illustration of the particle content of the standard model.	7
1.2	Measurements of the strong coupling constant, α_s , shown as a function of the energy scale, Q . The figure is taken from Ref. [3].	8
1.3	Illustration of a hard scattering process in a hadron-hadron collision.	13
1.4	Comparison between the predicted value for the cross section of various SM processes with the value obtained from measurements performed at the CMS experiment at the CERN LHC. The figure is taken from Ref. [8].	15
2.1	Illustration of the CERN accelerator complex. The figure is taken from Ref. [24].	19
2.2	Coordinate system used at the CMS experiment.	20
2.3	Illustration representing the CMS detector. Human figures are drawn to highlight the size of the detector. The figure is taken from Ref. [21].	20
2.4	Illustration of a slice of the CMS detector. The interaction of several SM particles with the different subdetectors is represented by the different coloured lines. The figure is taken from Ref. [25].	21
2.5	Sketch of a supercluster.	29
2.6	Average noise level, σ_{noise} , predicted in the ECAL barrel (<i>left</i>) and endcaps (<i>right</i>) for the Run 3 of data taking as a function of $ \eta $. The different coloured lines correspond to different total accumulated luminosities. The figures are taken from Ref. [46].	30
2.7	Illustration of the main steps of the ECAL clustering algorithm.	31
2.8	Energy resolution (<i>left</i>), signal efficiency (<i>middle</i>), and noise rate (<i>right</i>) as functions of the multipliers of σ_{noise} for the crystal and seeding thresholds in a representative energy and $ \eta $ bin.	32
2.9	Energy resolution (<i>upper left</i>), signal efficiency (<i>upper right</i>), and noise rate (<i>lower</i>) obtained with the tuned and reference thresholds as a function of $ \eta $. The performance is shown for simulated photons with transverse energy between 5 and 10 GeV that interact with the ECAL with conditions predicted for a total integrated luminosity of 450 fb^{-1}	33
3.1	Evolution of the inverse gauge couplings associated with the U(1), SU(2), and SU(3) gauge symmetries, for the case in which the sparticles have masses in the TeV range. The figure is taken from Ref. [1].	40
3.2	Feynman diagrams for gluino (<i>upper</i>) and squark (<i>lower</i>) production via gluon-gluon fusion (<i>left</i>) and quark-quark annihilation (<i>right</i>).	41

- 3.3 Cross section as a function of the sparticle mass, M , for gluino, light-flavour squark, and heavy-flavour squark pair-production in proton-proton collisions at $\sqrt{s} = 13$ TeV, computed at NNLO_{Approx}+NNLL accuracy. The width of the curves represents the associated theoretical uncertainties. The light-flavour squarks, $\tilde{u}$, $\tilde{d}$, $\tilde{s}$, and $\tilde{c}$, are assumed to be mass degenerate and the associated cross section is the sum between the left- and right-handed chiral states. 42
- 3.4 Feynman diagrams representing the gluino-mediated processes constrained in this work. The gluino decays into a light flavour (*left*), beauty (*middle*), and top (*right*) quarks are shown. 43
- 3.5 Feynman diagrams representing the squark-mediated processes constrained in this work. The direct production of light flavour (*left*), beauty (*middle*), and top (*right*) squark pairs is shown. 43
- 4.1 Schematic representation of events showing the SM decay $pp \rightarrow W^+$; $W^+ \rightarrow e^+ \nu_e$ (*left*), and the R -parity conserving SUSY decay $pp \rightarrow \tilde{q}\tilde{q}$; $\tilde{q} \rightarrow q\tilde{\chi}_1^0$ (*right*). The red and green colours highlight the visible and invisible final states, respectively. 48
- 4.2 Distribution of the M_{T2} variable shown for the simulated SM backgrounds (stacked histograms) and a representative signal (coloured line). The signal represents the pair-production of gluinos of mass $M = 1500$ GeV that decay into two beauty quarks and the lightest neutralino of mass $m_{\tilde{\chi}_1^0} = 100$ GeV. The normalisation of the signal is increased by a factor 10 to improve the readability. The figure was provided by UCSB. 49
- 4.3 Illustration of the binning in N_j and N_b for multijet events in the very-low and low H_T regions (*left*), and for the medium, high, and very-high H_T regions (*right*). Each blue square represents a bin in N_j and N_b 56
- 4.4 Distributions of M_{T2} (*upper left*), H_T (*upper right*), N_j (*lower left*), and N_b (*lower right*) for simulated background and signal events, shown for multijet events. Six signal models involving gluino-mediated decays into light-flavour quarks (cyan), beauty quarks (magenta), and top quarks (orange) are shown for non-compressed (solid lines) and compressed (dashed lines) mass scenarios. The normalisation of the signal is increased by a factor 10 for better readability. The plots were produced by the UCSB team. 60
- 4.5 Data to MC simulation comparison of the distribution of the kinematic variables H_T (*upper left*), N_j (*upper right*), N_b (*lower left*), and p_T^{jet} (*lower right*) for events in the same-flavour dilepton control region. The first three distributions are shown for events with at least two jets and the last one is shown for monojet events. All processes with a top quark are summed and the MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio between the data and MC yields is shown in the lower panel. 63

-
- 4.6 Data to MC simulation comparison of the M_{T2} distribution in the very-low H_T (left) and medium H_T (right) regions, for events in the same-flavour dilepton CR with at least two jets. All processes with a top quark are summed and the MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio between the data and MC yields is shown in the lower panel. The level of agreement between data and MC is similar for the other H_T regions. 64
- 4.7 $R^{\text{SF/DF}}$ factor as a function of the kinematic variables M_{T2} (upper left), H_T (upper right), N_j (lower left), and N_b (lower right). The measured central value $R^{\text{SF/DF}} = 1.06$ is shown as a solid black line, while the statistical and systematic uncertainties are represented with the black and red dashed lines, respectively. 65
- 4.8 Sketch illustrating the M_{T2} distribution in the same-flavour dilepton control region in a given topological region Ω . The main elements of the hybrid method used to extrapolate the M_{T2} shape to the $Z \rightarrow \nu\bar{\nu}$ estimate are reported. 66
- 4.9 Shape of the M_{T2} distribution for $Z \rightarrow \nu\bar{\nu}$ (left) and $Z \rightarrow \ell^+\ell^-$ (right) simulated events for different N_b regions. The comparison is shown for a representative region with $450 < H_T < 575$ GeV and $2 \leq N_j \leq 3$. The lower panel shows the ratio of the $N_b = 1$ and $N_b = 2$ entries to the $N_b = 0$ entries. 67
- 4.10 Comparison of the M_{T2} shape from $Z \rightarrow \ell^+\ell^-$ data and $Z \rightarrow \nu\bar{\nu}$ MC for events in the very-low H_T (left) and medium H_T (right) regions with $2 \leq N_j \leq 3$ and $N_b = 0$. The errors are statistical. 67
- 4.11 Comparison of data and MC distributions of H_T (upper left), N_j (upper right), and N_b (lower left), for events with at least two jets, and of p_T^{jet} (lower right), for events with exactly one jet. The distributions are shown for events in the single-lepton control region. All processes with a top quark are summed. The MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio of data to MC distributions is shown in the lower panel. . . 70
- 4.12 Comparison of data and MC distributions of the M_{T2} variable for events in the single-lepton control region in the very-low H_T (left) and medium H_T (right) regions. All processes with a top quark are summed. The MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio of data to MC distributions is shown in the lower panel. 71
- 4.13 Comparison of data and MC distributions of the H_T variable for multijet events in the single-lepton control region with $N_b = 0$ (left) and $N_b \geq 2$ (right). All processes with a top quark are summed. The MC distributions are normalised to data to allow for a shape comparison. The hashed bands represent the statistical uncertainty in MC. The last bin shows the overflow. The ratio of data to MC distributions is shown in the lower panel. 71

- 4.14 Comparison of the estimated background to the observed data events in each bin of N_b and p_T^{jet} in the monojet region (*upper*), and in each topological region Ω and M_{T2} bin in the very-low H_T region (*lower*). The hatched bands represent the full uncertainty in the background prediction. The labels “j” and “b” are short for N_j and N_b , respectively, and the p_T^{jet} and M_{T2} bins given on the x -axis of the upper and lower plots, respectively, are given in units of GeV. The lower panel shows the ratio of data to the background prediction. The multijet estimates (yellow) were provided by UCSB. 74
- 4.15 Comparison of the estimated background to the observed data events in each topological region Ω and M_{T2} bin in the low H_T (*upper*) and medium H_T (*lower*) regions. The hatched bands represent the full uncertainty in the background prediction. The labels “j” and “b” are short for N_j and N_b , respectively, and the M_{T2} bins given on the x -axis are given in units of GeV. The lower panel shows the ratio of data to the background prediction. The multijet estimates (yellow) were provided by UCSB. 75
- 4.16 Comparison of the estimated background to the observed data events in each topological region Ω and M_{T2} bin in the high H_T (*upper*) and very-high H_T (*lower*) regions. The hatched bands represent the full uncertainty in the background prediction. The labels “j” and “b” are short for N_j and N_b , respectively, and the M_{T2} bins given on the x -axis are given in units of GeV. The lower panel shows the ratio of data to the background prediction. The multijet estimates (yellow) were provided by UCSB. 76
- 4.17 Pull distribution obtained from the comparison of the observed data with the background prediction in each of the 282 search regions. The red curve shows the result from a fit with a gaussian function with mean $\mu = 0.04$ and standard deviation $\sigma = 0.8$ 77
- 4.18 The 95% CL upper exclusion limits for the gluino-mediated models, where the gluinos decay into light quarks (*upper left*), beauty quarks (*upper right*), or top quarks (*lower*), shown as a function of the masses of the gluino and the LSP. The thick black curve corresponds to the observed limit, while the thin black ones show the theoretical uncertainties in the signal cross section. The thick red curve corresponds to the median expected limits, while the thin red ones show the corresponding $\pm 1\sigma$ and $\pm 2\sigma$ ranges, where σ is the standard deviation. The coloured axis shows the observed upper limit on the signal cross section. The Feynman diagrams of the corresponding models are shown in Fig. 3.4. 79

4.19	The 95% CL upper exclusion limits for the squark-mediated models with pair-production of light-flavour squarks (<i>upper left</i>), beauty squarks (<i>upper right</i>), or top squarks (<i>lower</i>), shown as a function of the masses of the corresponding squark and the LSP. The thick black curve corresponds to the observed limit, while the thin black ones show the theoretical uncertainties in the signal cross section. The thick red curve corresponds to the median expected limits, while the thin red ones show the corresponding $\pm 1\sigma$ and $\pm 2\sigma$ ranges, where σ is the standard deviation. The coloured axis shows the observed upper limit on the signal cross section. In the lower figure, the grey dashed line shows the scenario in which the mass of the top squark equals the sum of the mass of the top quark and the mass of the LSP. The Feynman diagrams of the corresponding models are shown in Fig. 3.5.	80
4.20	Comparison of the limits from different searches targeting SUSY processes involving the pair-production of beauty squarks. The limits are shown from several searches performed at the CMS (<i>left</i>) and ATLAS (<i>right</i>) Collaborations. The corresponding references are provided in the legend of the figures. In particular, the limits from the present analysis and from the CMS analysis based on H_T^{miss} [81] are shown in red and blue in the left figure, respectively. The limits from the ATLAS search [82] targeting the same process as in the present analysis are shown in dark blue in the right figure. The left figure is taken from Ref. [83] and the right one from Ref. [84].	82
5.1	Illustration of the particle content of the ν MSM. The heavy neutrinos, N_1 , N_2 , and N_3 , are placed in arbitrary lepton families.	89
5.2	Exclusion limits on $ V_N ^2$ as a function of m_N , for the case in which the HNL mixes exclusively with the muon sector. The filled areas correspond to existing limits, while the dashed lines correspond to projections from future experiments. The figure is taken from Ref. [33].	92
6.1	Feynman diagrams representing the semileptonic (<i>upper</i>) and leptonic (<i>lower</i>) B meson decay into a charged lepton, ℓ_B , a hadronic system, X, in case of the semileptonic decay, and a neutrino, ν_{ℓ_B} , which contains a small admixture of a heavy neutrino, N. The N mass eigenstate propagates, and according to its admixture of the neutrino flavour eigenstate, ν_ℓ , decays into a charged lepton, $\ell^\pm$, and charged pion, $\pi^\mp$	94
6.2	B-parking rates of the L1 trigger (<i>left</i>) and HLT (<i>right</i>), and number of pileup vertices as a function of time for an arbitrary LHC fill. The modification of the trigger thresholds are indicated with the vertical lines. For the HLT, the black markers show the rate for the promptly reconstructed stream, while the blue ones show the rate for the B-parking stream. The figures are taken from Ref. [34]. . .	98

6.3	Branching fractions of the leptonic process $B_u \rightarrow eN$, the semileptonic process into a pseudoscalar meson $B_u \rightarrow eN\rho^0$, and the semileptonic process into a vector meson $B_u \rightarrow eN\rho^0$ shown as a function of the mass m_N . The branching fractions correspond to the scenario $(r_e, r_\mu, r_\tau) = (1, 0, 0)$ and $ V_N ^2 = 1$. The coloured curves are obtained with the tool implemented in this work and the black markers are points extracted from the results presented in Ref. [95].	100
6.4	(Left) Branching fraction of the leptonic and semileptonic decays divided by the total branching fraction, shown as a function of m_N . The fractions are computed for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ and $ V_N ^2 = 1$. (Right) Branching fractions as a function of m_N for $B_q \rightarrow \mu N X$ decays, with $q = (u, d, s, c)$, multiplied by the corresponding fragmentation fraction, f_q , listed in Table 6.3. The results are shown for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ and $ V_N ^2 = 1$	101
6.5	Branching fraction as a function of m_N for a heavy neutrino decaying via the hadronic, leptonic, and invisible channels. The contributions from the charged-current-mediated and neutral-current-mediated decays are summed. The widths are computed for the mixing scenario $(r_e, r_\mu, r_\tau) = (1/3, 1/3, 1/3)$. The coloured curves are obtained with the tool implemented in this work and the black markers are points extracted from the results presented in Ref. [95].	102
6.6	The upper panel shows the reduced HNL decay widths as a function of m_N . The indigo, red, and orange coloured solid curves show the widths obtained for the case in which the HNL decay involves a charged or neutral electron, muon, or tau lepton, respectively. The black dashed curve corresponds to the width of the $N \rightarrow \mu\pi$ decay, computed for the case where the HNL mixes exclusively with the muon sector. All the widths are shown for the Majorana scenario, and are hence a factor two larger than the widths for the Dirac-like scenario. The lower panel shows as a function of m_N the branching fraction of the $N \rightarrow \mu\pi$ decay, for the case in which the HNL mixes exclusively with the muon sector.	103
6.7	Distribution of the primary lepton p_T (upper left), displaced lepton p_T (upper right), invariant mass of the $\ell_B \ell^\pm \pi^\mp$ system (lower left), and L_{xy} of the DV (lower right). The distributions are shown for generator-level events from the (B_u, B_d, B_s) samples simulated with (i) $m_N = 1$ GeV, $c\tau_N = 1000$ mm, (ii) $m_N = 3$ GeV, $c\tau_N = 100$ mm, and (iii) $m_N = 4.5$ GeV, $c\tau_N = 1$ mm. No corrections are applied on the simulated events. The distributions are normalised to unit area.	105
6.8	Efficiency of the $\ell_B \ell^\pm \pi^\mp$ candidate reconstruction shown in bins of L_{xy} and $p_T(\mu^\pm \pi^\mp)$ for events in the dimuon channel simulated with $m_N = 3$ GeV and $c\tau_N = 2$ m.	107
6.9	Distributions of the displaced pion p_T (left) and $ d_{xy} /\sigma_{d_{xy}}$ (right) in data and in simulated events from the (B_u, B_d, B_s) samples corresponding to the three benchmark signal hypotheses. The distributions are shown for events in the dimuon channel and are normalised to unit area. The last bin shows the overflow. The corresponding preselection conditions are $p_T(\pi) > 1$ GeV and $ d_{xy} /\sigma_{d_{xy}}(\pi) > 5$, respectively.	109

6.10	Efficiency and fake rate of the <i>loose</i> identification criteria for the displaced muon shown for the three benchmark signal hypotheses as a function of the displaced muon p_T (<i>left</i>) and the DV L_{xy} (<i>right</i>).	110
6.11	Distribution of $L_{xy}/\sigma_{L_{xy}}$ (<i>left</i>) for data and the three signal benchmark hypotheses and distribution of the $\mu_B\mu^\pm\pi^\mp$ invariant mass (<i>right</i>) for (B_u , B_d , B_s) and B_c signals with $m_N = 4.5$ GeV and $c\tau_N = 1$ mm. The distributions are shown for events in the dimuon channel after the preselection and are normalised to unit area. The boundaries of the analysis categories are indicated with the orange lines. The last bin shows the overflow.	113
6.12	The distributions of the pion p_T (<i>upper left</i>), $m(\mu_B\mu^\pm\pi^\mp)$ (<i>upper right</i>), $\cos\theta$ (<i>lower left</i>), and pion $ d_{xy} /\sigma_{d_{xy}}$ (<i>lower right</i>) are shown for data and for a signal hypothesis of $m_N = 2$ GeV and $c\tau_N = 100$ mm. The data are selected in the mass window of size $\pm 10\sigma$ around $m_N = 2$ GeV and correspond to an integrated luminosity of 5.2 fb^{-1} . The distributions, which are normalised to unit area, are shown in the category with high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ in the dimuon channel.	117
6.13	Shape of the distribution of the pNN score (<i>left</i>) and ROC curve (<i>right</i>) for signal hypotheses with $m_N = 1.0$ GeV and $c\tau_N = 10.0, 100.0$, and 1000.0 mm for the pNN trained in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel.	118
6.14	(<i>Left</i>) Performance of the pNN assessed with the AUC as a function of m_N for events in the $50 < L_{xy}/\sigma_{L_{xy}} < 150$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. The AUC at unity corresponds to a perfect separation between signal and background. The different coloured curves show different $c\tau_N$ hypotheses. (<i>Right</i>) Comparison of the AUC for a pNN trained on masses $m_N = 1.0, 1.5, 2.0, 3.0$, and 4.5 GeV (blue) and for a NN trained on mass $m_N = 2$ GeV (red), for events in the $50 < L_{xy}/\sigma_{L_{xy}} < 150$, OS and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. The points have $c\tau_N = 10$ mm. The full circles show mass points on which the pNN and NN were trained on, while the open circles show mass points that were not trained on.	119
6.15	Distribution of the p_T of the triggering μ_B shown without (<i>left</i>) and with (<i>right</i>) the application of the corrections for the trigger efficiency. The distributions are compared between data and simulated soft QCD events in the dimuon channel with a non-zero charge of the $\mu^\pm\pi^\mp$ system.	121
6.16	Ratio of the trigger efficiencies in data and MC in bins of the probe muon p_T and $ d_{xy} /\sigma_{d_{xy}}$. The uncertainty corresponds to the propagated uncertainty in the number of events returned by the fits to the dimuon mass spectra in data and simulation.	122
6.17	Distribution of the $K^\pm\mu^+\mu^-$ invariant mass. The signal peak is observed at the mass of the B_u meson. The orange, green, and red curves show the contributions from signal, composite background, and combinatorial background, respectively, while the blue curve shows the fit to signal plus background.	128

6.18	Distribution of $m(\ell^\pm\pi^\mp)$ for simulated events with $m_N = 1$ GeV and $c\tau_N = 1000$ mm in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. The blue curve corresponds to the fit to the data using a double-sided Crystal Ball function.	129
6.19	Resolution as a function of the signal mass m_N for events in the dimuon channel. For each mass point, the resolution is averaged over the given $c\tau_N$ hypotheses. The points are fitted with a polynomial function of first order.	130
6.20	Distribution of the $\mu^\pm\pi^\mp$ invariant mass (<i>left</i>) and $L_{xy}/\sigma_{L_{xy}}$ (<i>right</i>) for simulated QCD background events in the dimuon channel. The contributions from B cascade decays, combinatorial background, and instrumental background are indicated with the magenta, green, and grey colour, respectively.	131
6.21	Distribution of $m(\ell^\pm\pi^\mp)$ in the mass window around 2 GeV in the low $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. (<i>Left</i>) The coloured dashed lines represent the different functions of the set $\mathcal{S}$ obtained in the side-bands. The grey hashed region represents the signal region of size $\pm 3\sigma$. (<i>Right</i>) The red curve shows the result of the background-only fit to the data and the green curve shows an example of the expected Majorana signal with $m_N = 2$ GeV and $c\tau_N = 10$ mm for the mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$	133
6.22	Distribution of score of the pNN evaluated in the $B^\pm \rightarrow J/\psi(\rightarrow \mu^+\mu^-)K^\pm$ channel for data and MC events in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category in the dimuon channel. The pNN is evaluated with a mass key equal to 4.5 GeV.	136
6.23	Fits to the $m(\ell^\pm\pi^\mp)$ distributions for signals with mass $m_N = 1.0$ GeV (<i>upper left</i>), 1.5 GeV (<i>upper right</i>), 2.0 GeV (<i>lower left</i>), and 2.5 GeV (<i>lower right</i>), in the high $L_{xy}/\sigma_{L_{xy}}$, OS, and low $m(\ell_B\ell^\pm\pi^\mp)$ category of the dimuon channel. The signal hypotheses correspond to $c\tau_N = 10000, 2000, 700$, and 100 mm, respectively. The blue curve shows the signal-plus-background fit, while the green and red curves show the individual signal and background components, respectively. The yellow band corresponds to the total systematic plus statistical uncertainty.	138
6.24	Expected and observed 95% CL limits on $ V_N ^2$ as a function of m_N in the Majorana (<i>left</i>) and Dirac-like (<i>right</i>) scenarios. The limits are shown for the muon-exclusive mixing scenario $(r_e, r_\mu, r_\tau) = (0, 1, 0)$. Results from the CMS [107, 113], ATLAS [105], LHCb [103], and Belle [102] Collaborations are reported. The grey bands indicate the mass regions that are vetoed, as listed in Table 6.6.	139
6.25	Expected and observed 95% CL limits on $ V_N ^2$ as a function of m_N in the Majorana (<i>left</i>) and Dirac-like (<i>right</i>) scenarios. The limits are shown for the mixed-flavour mixing scenarios $(r_e, r_\mu, r_\tau) = (0, 1/2, 1/2)$ (<i>upper</i>), $(r_e, r_\mu, r_\tau) = (1/2, 1/2, 0)$ (<i>middle</i>), and $(r_e, r_\mu, r_\tau) = (1/3, 1/3, 1/3)$ (<i>lower</i>). Results from the CMS [112] Collaboration are reported. The grey bands indicate the mass regions that are vetoed, as listed in Table 6.6.	141

6.26	Observed 95% CL upper limits on $ V_N ^2$ as functions of the mixing ratios (r_e , r_μ , r_τ) for fixed HNL masses of 1 GeV (<i>upper</i>), 1.5 GeV (<i>middle</i>), and 2 GeV (<i>lower</i>), in the Majorana (<i>left</i>) and Dirac-like (<i>right</i>) scenarios. The range of values on the z -axis is not kept fixed across the different plots. The red crosses indicate that there is no exclusion for that point. The orientation of the value markers on each axis indicates the direction of reading.	143
6.27	Observed 95% CL lower limits on $c\tau_N$ as functions of the mixing ratios (r_e , r_μ , r_τ) for fixed HNL masses of 1 GeV (<i>upper</i>), 1.5 GeV (<i>middle</i>), and 2 GeV (<i>lower</i>), in the Majorana (<i>left</i>) and Dirac-like (<i>right</i>) scenarios. The range of values on the z -axis is kept fixed across the different plots. The red crosses indicate that there is no exclusion for that point. The orientation of the value markers on each axis indicates the direction of reading.	144
A.1	Relative difference of the resolution obtained with the tuned and reference energy thresholds, shown in different bins of the simulated photon energy and $ \eta $. The performance is shown for simulated photons interacting with the ECAL with conditions corresponding to an integrated luminosity of 450 fb^{-1} . The red curve is a reference corresponding to a transverse energy of 5 GeV. The blank bins with a hash symbol are not populated by clusters as the energy is below the energy thresholds.	150
A.2	Relative difference of the efficiency obtained with the tuned and reference energy thresholds, shown in different bins of the simulated photon energy and $ \eta $. The performance is shown for simulated photons interacting with the ECAL with conditions corresponding to an integrated luminosity of 450 fb^{-1} . The red curve is a reference corresponding to a transverse energy of 5 GeV. The blank bins with a hash symbol are not populated by clusters as the energy is below the energy thresholds.	151
A.3	Relative difference of the noise rate obtained with the tuned and reference energy thresholds, shown in different bins of the cluster energy and $ \eta $. The performance is shown for simulated photons interacting with the ECAL with conditions corresponding to an integrated luminosity of 450 fb^{-1} . The red curve is a reference corresponding to a transverse energy of 5 GeV. The blank bins with a hash symbol are not populated by clusters as the energy is below the energy thresholds. The dash symbol indicates bins in which the noise rate is null with the reference thresholds.	152

List of tables

2.1	List of runs of data taking that have taken place up to the moment of writing. For each run, the time lapse, the maximal centre-of-mass energy, as well as the delivered integrated luminosity are reported. The asterisk (*) symbol indicates that the values are reported as of the time of writing: the Run 3 of data taking is expected to last until 2026, and reach an integrated luminosity of about 160 fb^{-1} .	18
3.1	Content of sparticles in the MSSM.	38
3.2	Neutralinos and charginos in the MSSM.	39
3.3	List of the three gluino-mediated and three squark-mediated SUSY models considered in this work. For each signal model, the sparticle production and decay modes are reported.	44
4.1	List of previous iterations of the search together with the reference to their respective publication. The searches were performed using proton-proton collision data collected at $\sqrt{s} = 7, 8, \text{ or } 13 \text{ TeV}$ and corresponding to different integrated luminosities.	46
4.2	List of integrated luminosity per year of data taking.	50
4.3	List of hadronic, invisible, and mixed triggers used during the three years of data taking. For each trigger, the HLT thresholds on the value of H_T , p_T of any jet, p_T^{miss} , and H_T^{miss} are reported. The dash symbol “–” indicates that there is no threshold on the corresponding quantity.	51
4.4	List of simulated background and signal samples. For each sample, the frameworks used for the event generation and the simulation of the detector response are reported.	52
4.5	Jet identification criteria in monojet ($N_j = 1$) and multijet ($N_j > 1$) events for the different data sets. The dash symbol “–” indicates that no condition on the corresponding quantity is applied.	53
4.6	List of final state objects and corresponding search regions that trigger the additional HCAL veto. Events in the 2018 data sets impacted by the two HCAL modules failure are rejected if they contain at least one of these objects in the corresponding search region and with $-4.7 < \eta < -1.4$ and $-1.6 < \phi < -0.8$.	55
4.7	Definition of the signal regions for the very-low, low, and medium H_T multijet regions. The labels “j” and “b” are short for N_j and N_b , respectively.	57
4.8	Definition of the signal regions for the high and very-high H_T multijet regions. The labels “j” and “b” are short for N_j and N_b , respectively.	58

4.9	Definition of the signal regions for the monojet regions. The labels “j” and “b” are short for N_j and N_b , respectively.	58
4.10	Sources of systematic uncertainties in the signal yield, together with their typical magnitude ranges across the search categories.	61
4.11	Sources of systematic uncertainties in the $Z \rightarrow \nu\bar{\nu}$ background estimate, together with their typical magnitude ranges across the search categories.	68
4.12	Sources of systematic uncertainties in the lost-lepton background estimate, together with their typical magnitude ranges across the search categories.	73
4.13	Highest 95% CL exclusion limits on directly produced sparticle mass, M , and on the LSP mass, $m_{\tilde{\chi}_1^0}$, obtained for the different simplified SUSY models. The limits are presented from this analysis (using 137 fb^{-1}) and from the previous iteration of this analysis [66] (using 36 fb^{-1}). The relative difference between the two sets of results is also reported.	81
6.1	List of the HLT trigger conditions. Each line represents the combination of the minimal thresholds applied on the muon p_T and $ d_{xy} /\sigma_{d_{xy}}$ for a given HLT path.	97
6.2	List of the main heavy neutrino production decays from the B_u , B_d , B_s , and B_c mesons. The decay modes are categorised into leptonic processes, semileptonic processes involving a pseudoscalar meson, and semileptonic processes involving a vector meson. The dash symbol “–” indicates that there is no corresponding decay mode.	99
6.3	Values of the fragmentation fractions f_q ($q = u, d, s, c$) for the different B meson species and their corresponding reference. The fragmentation fraction f_q is defined as the probability for a b quark to hadronise into a B_q meson.	100
6.4	List of charged-current-mediated and neutral-current-mediated heavy neutrino decays in the hadronic, leptonic, and invisible channels. The quarks u and d can have any of the up-type or down-type quark flavour, respectively, and the quark q can have any flavour. The indices α and β run over the three lepton flavours (e, μ , τ). The dash symbol “–” indicates that there is no corresponding decay mode. In the signal process, the charged-current-mediated hadronic decay into $\ell\pi$ is targeted.	101
6.5	List of the event categories.	112
6.6	List of SM resonances and corresponding vetoes in the $\mu_B\mu^\pm$ and $\mu_B\pi^\mp$ invariant mass spectra (<i>upper part</i>) and in the $\mu^\pm\pi^\mp$ invariant mass spectrum (<i>lower part</i>). For each process, the categories in which the veto is applied is reported and the presence of misidentified particles is indicated.	132
6.7	List of systematic uncertainties impacting the expected signal event yield. The ranges in the value represent the uncertainties across the different event categories.	134
6.8	Most stringent observed upper limits on $ V_N ^2$ at 95% CL for the different mixing scenarios. The minimum excluded value of $ V_N ^2$ is reported together with the signal mass at which it is obtained.	142

6.9	Summary of the most stringent upper limits at 95% CL on $ V_N ^2$, obtained for the muon-exclusive mixing scenario. The minimum excluded value of $ V_N ^2$ is reported for masses $m_N = 1.0, 1.5$ and 2.0 , and for the Majorana and Dirac-like scenarios.	145
6.10	Summary of the most stringent lower observed limits at 95% CL on $c\tau_N$, obtained for the muon-exclusive mixing scenario. The maximum excluded value of $c\tau_N$ is reported for masses $m_N = 1.0, 1.5$ and 2.0 , and for the Majorana and Dirac-like scenarios.	145

Bibliography

- [1] M. Thomson, “Modern particle physics”. Cambridge University Press, New York, (2013). [doi:10.1017/CB09781139525367](https://doi.org/10.1017/CB09781139525367), ISBN 978-1-107-03426-6.
- [2] A. Rubbia, “Phenomenology of Particle Physics”. Cambridge University Press, (2022).
- [3] Particle Data Group Collaboration, “Review of Particle Physics”, *PTEP* **2022** (2022) 083C01, [doi:10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [4] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons”, *Phys. Rev. Lett.* **13** (1964) 321, [doi:10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321).
- [5] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons”, *Phys. Rev. Lett.* **13** (1964) 508, [doi:10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [6] M. Goldhaber, L. Grodzins, and A. W. Sunyar, “Helicity of Neutrinos”, *Phys. Rev.* **109** (1958) 1015, [doi:10.1103/PhysRev.109.1015](https://doi.org/10.1103/PhysRev.109.1015).
- [7] M. C. Gonzalez-Garcia and M. Maltoni, “Phenomenology with Massive Neutrinos”, *Phys. Rept.* **460** (2008) 1, [doi:10.1016/j.physrep.2007.12.004](https://doi.org/10.1016/j.physrep.2007.12.004), [arXiv:0704.1800](https://arxiv.org/abs/0704.1800).
- [8] CMS Collaboration, “Stairway to discovery: a report on the CMS programme of cross section measurements from millibarns to femtobarns”, [arXiv:2405.18661](https://arxiv.org/abs/2405.18661). (2024).
- [9] UA1 Collaboration, “Experimental Observation of Lepton Pairs of Invariant Mass Around $95 \text{ GeV}/c^2$ at the CERN SPS Collider”, *Phys. Lett. B* **126** (1983) 398, [doi:10.1016/0370-2693\(83\)90188-0](https://doi.org/10.1016/0370-2693(83)90188-0).
- [10] UA2 Collaboration, “Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $\bar{p}p$ Collider”, *Phys. Lett. B* **129** (1983) 130, [doi:10.1016/0370-2693\(83\)90744-X](https://doi.org/10.1016/0370-2693(83)90744-X).
- [11] CMS Collaboration, “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”, *Phys. Lett. B* **716** (2012) 30, [doi:10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021), [arXiv:1207.7235](https://arxiv.org/abs/1207.7235).
- [12] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Phys. Lett. B* **716** (2012) 1, [doi:10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020), [arXiv:1207.7214](https://arxiv.org/abs/1207.7214).

- [13] S. Weinberg, “Implications of Dynamical Symmetry Breaking”, *Phys. Rev. D* **13** (1976) 974, [doi:10.1103/PhysRevD.19.1277](https://doi.org/10.1103/PhysRevD.19.1277). [Addendum: *Phys.Rev.D* 19, 1277 (1979)].
- [14] S. Bilenky, “Neutrino oscillations: From a historical perspective to the present status”, *Nucl. Phys. B* **908** (2016) 2, [doi:10.1016/j.nuclphysb.2016.01.025](https://doi.org/10.1016/j.nuclphysb.2016.01.025), [arXiv:1602.00170](https://arxiv.org/abs/1602.00170).
- [15] G. R. Farrar and M. E. Shaposhnikov, “Baryon asymmetry of the universe in the standard electroweak theory”, *Phys. Rev. D* **50** (1994) 774, [doi:10.1103/PhysRevD.50.774](https://doi.org/10.1103/PhysRevD.50.774), [arXiv:hep-ph/9305275](https://arxiv.org/abs/hep-ph/9305275).
- [16] J. Silk et al., “Particle dark matter: observations, models and searches”. Cambridge Univ. Press, Cambridge, (2010). [doi:10.1017/CB09780511770739](https://doi.org/10.1017/CB09780511770739), ISBN 978-1-107-65392-4.
- [17] O. S. Bruning et al., “LHC Design Report”, (2004). [doi:10.5170/CERN-2004-003-V-1](https://doi.org/10.5170/CERN-2004-003-V-1).
- [18] “Design Report Tevatron 1 project”, (1984). <https://cds.cern.ch/record/1478620>.
- [19] I. Zurbano Fernandez et al., “High-Luminosity Large Hadron Collider (HL-LHC): Technical design report”, 12, (2020). [doi:10.23731/CYRM-2020-0010](https://doi.org/10.23731/CYRM-2020-0010).
- [20] ATLAS Collaboration, “The ATLAS Experiment at the CERN Large Hadron Collider”, *JINST* **3** (2008) S08003, [doi:10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003).
- [21] CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **3** (2008) S08004, [doi:10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004).
- [22] LHCb Collaboration, “The LHCb Detector at the LHC”, *JINST* **3** (2008) S08005, [doi:10.1088/1748-0221/3/08/S08005](https://doi.org/10.1088/1748-0221/3/08/S08005).
- [23] ALICE Collaboration, “The ALICE experiment at the CERN LHC”, *JINST* **3** (2008) S08002, [doi:10.1088/1748-0221/3/08/S08002](https://doi.org/10.1088/1748-0221/3/08/S08002).
- [24] F. Marcastel, “CERN’s Accelerator Complex. La chaîne des accélérateurs du CERN”, (2013). <https://cds.cern.ch/record/1621583>.
- [25] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017) P10003, [doi:10.1088/1748-0221/12/10/P10003](https://doi.org/10.1088/1748-0221/12/10/P10003), [arXiv:1706.04965](https://arxiv.org/abs/1706.04965).
- [26] CMS Collaboration, “The CMS tracker system project: Technical Design Report”, (1997).
- [27] CMS Collaboration, “CMS Technical Design Report for the Pixel Detector Upgrade”, [doi:10.2172/1151650](https://doi.org/10.2172/1151650). (2012).
- [28] CMS Collaboration, “The CMS electromagnetic calorimeter project: Technical Design Report”, (1997).

-
- [29] CMS Collaboration, “The CMS hadron calorimeter project: Technical Design Report”, (1997).
 - [30] CMS Collaboration, “The CMS muon project: Technical Design Report”, (1997).
 - [31] CMS Collaboration, “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020) P10017, [doi:10.1088/1748-0221/15/10/P10017](https://doi.org/10.1088/1748-0221/15/10/P10017), [arXiv:2006.10165](https://arxiv.org/abs/2006.10165).
 - [32] CMS Collaboration, “The CMS trigger system”, *JINST* **12** (2017) P01020, [doi:10.1088/1748-0221/12/01/P01020](https://doi.org/10.1088/1748-0221/12/01/P01020), [arXiv:1609.02366](https://arxiv.org/abs/1609.02366).
 - [33] C. Antel et al., “Feebly-interacting particles: FIPs 2022 Workshop Report”, *Eur. Phys. J. C* **83** (2023) 1122, [doi:10.1140/epjc/s10052-023-12168-5](https://doi.org/10.1140/epjc/s10052-023-12168-5), [arXiv:2305.01715](https://arxiv.org/abs/2305.01715).
 - [34] CMS Collaboration, “Enriching the Physics Program of the CMS Experiment via Data Scouting and Data Parking”, [arXiv:2403.16134](https://arxiv.org/abs/2403.16134). (2024), Accepted by *Phys. Rept.*
 - [35] CMS Collaboration, “Recording and reconstructing 10 billion unbiased b hadron decays in CMS”, (2019). <https://cds.cern.ch/record/2704495>.
 - [36] CMS Collaboration, “Test of lepton flavor universality in $B^\pm \rightarrow K^\pm \mu^+ \mu^-$ and $B^\pm \rightarrow K^\pm e^+ e^-$ decays in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Rept. Prog. Phys.* **87** (2024) 077802, [doi:10.1088/1361-6633/ad4e65](https://doi.org/10.1088/1361-6633/ad4e65), [arXiv:2401.07090](https://arxiv.org/abs/2401.07090).
 - [37] T. Speer et al., “Track reconstruction in the CMS tracker”, *Nucl. Instrum. Meth. A* **559** (2006) 143, [doi:10.1016/j.nima.2005.11.207](https://doi.org/10.1016/j.nima.2005.11.207).
 - [38] R. Fruhwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Meth. A* **262** (1987) 444, [doi:10.1016/0168-9002\(87\)90887-4](https://doi.org/10.1016/0168-9002(87)90887-4).
 - [39] K. Rose, “Deterministic annealing for clustering, compression, classification, regression, and related optimization problems”, *IEEE Proc.* **86** (1998) 2210, [doi:10.1109/5.726788](https://doi.org/10.1109/5.726788).
 - [40] 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](https://doi.org/10.1088/1126-6708/2008/04/063), [arXiv:0802.1189](https://arxiv.org/abs/0802.1189).
 - [41] CMS Collaboration, “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”, *JINST* **12** (2017) P02014, [doi:10.1088/1748-0221/12/02/P02014](https://doi.org/10.1088/1748-0221/12/02/P02014), [arXiv:1607.03663](https://arxiv.org/abs/1607.03663).
 - [42] CMS Collaboration, “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV”, *JINST* **13** (2018) P05011, [doi:10.1088/1748-0221/13/05/P05011](https://doi.org/10.1088/1748-0221/13/05/P05011), [arXiv:1712.07158](https://arxiv.org/abs/1712.07158).
 - [43] CMS Collaboration, “Performance of CMS Muon Reconstruction in pp Collision Events at $\sqrt{s} = 7$ TeV”, *JINST* **7** (2012) P10002, [doi:10.1088/1748-0221/7/10/P10002](https://doi.org/10.1088/1748-0221/7/10/P10002), [arXiv:1206.4071](https://arxiv.org/abs/1206.4071).

- [44] CMS Collaboration, “Performance of Electron Reconstruction and Selection with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV”, *JINST* **10** (2015) P06005, [doi:10.1088/1748-0221/10/06/P06005](https://doi.org/10.1088/1748-0221/10/06/P06005), [arXiv:1502.02701](https://arxiv.org/abs/1502.02701).
- [45] CMS Collaboration, “Performance of Photon Reconstruction and Identification with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV”, *JINST* **10** (2015) P08010, [doi:10.1088/1748-0221/10/08/P08010](https://doi.org/10.1088/1748-0221/10/08/P08010), [arXiv:1502.02702](https://arxiv.org/abs/1502.02702).
- [46] “ECAL Clustering for Run 3”, (2022). <https://cds.cern.ch/record/2812783>.
- [47] M. J. Oreglia, “A study of the reactions $\psi' \rightarrow \gamma\gamma\psi$ ”, 1980. SLAC Report SLAC-R-236, <http://www.slac.stanford.edu/cgi-wrap/getdoc/slac-r-236.pdf>.
- [48] J. E. Gaiser, “Charmonium spectroscopy from radiative decays of the J/ψ and ψ' ”. PhD thesis, Stanford University, 1982. SLAC Report SLAC-R-255.
- [49] A.-M. Lyon, “Optimisation of the CMS ECAL clustering algorithms in view of LHC-Run 3”, *J. Phys. Conf. Ser.* **2374** (2022) 012015, [doi:10.1088/1742-6596/2374/1/012015](https://doi.org/10.1088/1742-6596/2374/1/012015).
- [50] S. P. Martin, “A Supersymmetry primer”, *Adv. Ser. Direct. High Energy Phys.* **18** (1998) 1, [doi:10.1142/9789812839657_0001](https://doi.org/10.1142/9789812839657_0001), [arXiv:hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356).
- [51] L. Pape and D. Treille, “Supersymmetry facing experiment: Much ado (already) about nothing (yet)”, *Rept. Prog. Phys.* **69** (2006) 2843, [doi:10.1088/0034-4885/69/11/R01](https://doi.org/10.1088/0034-4885/69/11/R01).
- [52] R. Barbier et al., “R-parity violating supersymmetry”, *Phys. Rept.* **420** (2005) 1, [doi:10.1016/j.physrep.2005.08.006](https://doi.org/10.1016/j.physrep.2005.08.006), [arXiv:hep-ph/0406039](https://arxiv.org/abs/hep-ph/0406039).
- [53] J. R. Ellis et al., “Supersymmetric Relics from the Big Bang”, *Nucl. Phys. B* **238** (1984) 453, [doi:10.1016/0550-3213\(84\)90461-9](https://doi.org/10.1016/0550-3213(84)90461-9).
- [54] E. Witten, “Dynamical Breaking of Supersymmetry”, *Nucl. Phys. B* **188** (1981) 513, [doi:10.1016/0550-3213\(81\)90006-7](https://doi.org/10.1016/0550-3213(81)90006-7).
- [55] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter”, *Phys. Rept.* **267** (1996) 195, [doi:10.1016/0370-1573\(95\)00058-5](https://doi.org/10.1016/0370-1573(95)00058-5), [arXiv:hep-ph/9506380](https://arxiv.org/abs/hep-ph/9506380).
- [56] J. R. Ellis, S. Kelley, and D. V. Nanopoulos, “Probing the desert using gauge coupling unification”, *Phys. Lett. B* **260** (1991) 131, [doi:10.1016/0370-2693\(91\)90980-5](https://doi.org/10.1016/0370-2693(91)90980-5).
- [57] J. R. Ellis, G. Ridolfi, and F. Zwirner, “Radiative corrections to the masses of supersymmetric Higgs bosons”, *Phys. Lett. B* **257** (1991) 83, [doi:10.1016/0370-2693\(91\)90863-L](https://doi.org/10.1016/0370-2693(91)90863-L).
- [58] Particle Data Group Collaboration, “Review of particle physics”, *J. Phys. G* **37** (2010) 075021, [doi:10.1088/0954-3899/37/7A/075021](https://doi.org/10.1088/0954-3899/37/7A/075021).

-
- [59] W. Beenakker et al., “NNLL-fast: predictions for coloured supersymmetric particle production at the LHC with threshold and Coulomb resummation”, *JHEP* **12** (2016) 133, [doi:10.1007/JHEP12\(2016\)133](#), [arXiv:1607.07741](#).
 - [60] B. Fuks, M. Klasen, D. R. Lamprea, and M. Rothering, “Precision predictions for electroweak superpartner production at hadron colliders with Resummino”, *Eur. Phys. J. C* **73** (2013) 2480, [doi:10.1140/epjc/s10052-013-2480-0](#), [arXiv:1304.0790](#).
 - [61] 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](#).
 - [62] 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](#).
 - [63] CMS Collaboration, “Search for supersymmetry in hadronic final states using M_{T2} in pp collisions at $\sqrt{s} = 7$ TeV”, *JHEP* **10** (2012) 018, [doi:10.1007/JHEP10\(2012\)018](#), [arXiv:1207.1798](#).
 - [64] CMS Collaboration, “Searches for supersymmetry using the M_{T2} variable in hadronic events produced in pp collisions at 8 TeV”, *JHEP* **05** (2015) 078, [doi:10.1007/JHEP05\(2015\)078](#), [arXiv:1502.04358](#).
 - [65] CMS Collaboration, “Search for new physics with the M_{T2} variable in all-jets final states produced in pp collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **10** (2016) 006, [doi:10.1007/JHEP10\(2016\)006](#), [arXiv:1603.04053](#).
 - [66] CMS Collaboration, “Search for new phenomena with the M_{T2} variable in the all-hadronic final state produced in proton–proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **77** (2017) 710, [doi:10.1140/epjc/s10052-017-5267-x](#), [arXiv:1705.04650](#).
 - [67] M. Masciovecchio, “Searches for strong SUSY at ATLAS and CMS at the LHC Run 2”, in *54th Rencontres de Moriond on Electroweak Interactions and Unified Theories*, p. 97. (2019).
 - [68] CMS Collaboration, “Searches for physics beyond the standard model with the M_{T2} variable in hadronic final states with and without disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **80** (2020) 3, [doi:10.1140/epjc/s10052-019-7493-x](#), [arXiv:1909.03460](#).
 - [69] UA1 Collaboration, “Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV”, *Phys. Lett. B* **122** (1983) 103, [doi:10.1016/0370-2693\(83\)91177-2](#).

- [70] UA2 Collaboration, “Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider”, *Phys. Lett. B* **122** (1983) 476, [doi:10.1016/0370-2693\(83\)91605-2](#).
- [71] C. G. Lester and D. J. Summers, “Measuring masses of semi-invisibly decaying particles pair produced at hadron colliders”, *Phys. Lett. B* **463** (1999) 99, [doi:10.1016/S0370-2693\(99\)00945-4](#), [arXiv:hep-ph/9906349](#).
- [72] CMS Collaboration, “CMS technical design report, volume II: Physics performance”, *J. Phys. G* **34** (2007) 995, [doi:10.1088/0954-3899/34/6/S01](#).
- [73] 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](#).
- [74] T. Sjöstrand et al., “An introduction to PYTHIA 8.2”, *Comput. Phys. Commun.* **191** (2015) 159, [doi:10.1016/j.cpc.2015.01.024](#), [arXiv:1410.3012](#).
- [75] GEANT4 Collaboration, “GEANT4—a simulation toolkit”, *Nucl. Instrum. Meth. A* **506** (2003) 250, [doi:10.1016/S0168-9002\(03\)01368-8](#).
- [76] A. Giammanco, “The Fast Simulation of the CMS Experiment”, *J. Phys. Conf. Ser.* **513** (2014) 022012, [doi:10.1088/1742-6596/513/2/022012](#).
- [77] CMS Collaboration, “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector”, *JINST* **14** (2019) P07004, [doi:10.1088/1748-0221/14/07/P07004](#), [arXiv:1903.06078](#).
- [78] CMS Collaboration, “Measurements of $t\bar{t}$ cross sections in association with b jets and inclusive jets and their ratio using dilepton final states in pp collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **776** (2018) 355, [doi:10.1016/j.physletb.2017.11.043](#), [arXiv:1705.10141](#).
- [79] 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** (2013) 2501].
- [80] A. L. Read, “Presentation of search results: The CL_s technique”, *J. Phys. G* **28** (2002) 2693, [doi:10.1088/0954-3899/28/10/313](#).
- [81] CMS Collaboration, “Search for supersymmetry in proton-proton collisions at 13 TeV in final states with jets and missing transverse momentum”, *JHEP* **10** (2019) 244, [doi:10.1007/JHEP10\(2019\)244](#), [arXiv:1908.04722](#).
- [82] ATLAS Collaboration, “Search for new phenomena in final states with b -jets and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”, *JHEP* **05** (2021) 093, [doi:10.1007/JHEP05\(2021\)093](#), [arXiv:2101.12527](#).

-
- [83] CMS Collaboration, “CMS SUSY Public Physics Results”, (2024).
<https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS>.
 - [84] ATLAS Collaboration, “SUSY March 2023 Summary Plot Update”,.
 - [85] Z. Liu and B. Tweedie, “The Fate of Long-Lived Superparticles with Hadronic Decays after LHC Run 1”, *JHEP* **06** (2015) 042, [doi:10.1007/JHEP06\(2015\)042](https://doi.org/10.1007/JHEP06(2015)042),
[arXiv:1503.05923](https://arxiv.org/abs/1503.05923).
 - [86] Super-Kamiokande Collaboration, “Evidence for oscillation of atmospheric neutrinos”, *Phys. Rev. Lett.* **81** (1998) 1562, [doi:10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562),
[arXiv:hep-ex/9807003](https://arxiv.org/abs/hep-ex/9807003).
 - [87] SNO Collaboration, “Measurement of the total active ^8B solar neutrino flux at the Sudbury Neutrino Observatory with enhanced neutral current sensitivity”, *Phys. Rev. Lett.* **92** (2004) 181301, [doi:10.1103/PhysRevLett.92.181301](https://doi.org/10.1103/PhysRevLett.92.181301),
[arXiv:nuc1-ex/0309004](https://arxiv.org/abs/nuc1-ex/0309004).
 - [88] Planck Collaboration, “Planck 2018 results. VI. Cosmological parameters”, *Astron. Astrophys.* **641** (2020) A6, [doi:10.1051/0004-6361/201833910](https://doi.org/10.1051/0004-6361/201833910), [arXiv:1807.06209](https://arxiv.org/abs/1807.06209).
[Erratum: *Astron. Astrophys.* 652, C4 (2021)].
 - [89] M. Drewes, “The Phenomenology of Right Handed Neutrinos”, *Int. J. Mod. Phys. E* **22** (2013) 1330019, [doi:10.1142/S0218301313300191](https://doi.org/10.1142/S0218301313300191), [arXiv:1303.6912](https://arxiv.org/abs/1303.6912).
 - [90] P. Minkowski, “ $\mu \rightarrow e\gamma$ at a rate of one out of 10^9 muon decays?”, *Phys. Lett. B* **67** (1977) 421, [doi:10.1016/0370-2693\(77\)90435-X](https://doi.org/10.1016/0370-2693(77)90435-X).
 - [91] A. M. Abdullahi et al., “The present and future status of heavy neutral leptons”, *J. Phys. G* **50** (2023) 020501, [doi:10.1088/1361-6471/ac98f9](https://doi.org/10.1088/1361-6471/ac98f9), [arXiv:2203.08039](https://arxiv.org/abs/2203.08039).
 - [92] K. N. Abazajian et al., “Light Sterile Neutrinos: A White Paper”, [arXiv:1204.5379](https://arxiv.org/abs/1204.5379). (2012).
 - [93] T. Asaka, S. Blanchet, and M. Shaposhnikov, “The nuMSM, dark matter and neutrino masses”, *Phys. Lett. B* **631** (2005) 151, [doi:10.1016/j.physletb.2005.09.070](https://doi.org/10.1016/j.physletb.2005.09.070),
[arXiv:hep-ph/0503065](https://arxiv.org/abs/hep-ph/0503065).
 - [94] T. Asaka and M. Shaposhnikov, “The ν MSM, dark matter and baryon asymmetry of the universe”, *Phys. Lett. B* **620** (2005) 17, [doi:10.1016/j.physletb.2005.06.020](https://doi.org/10.1016/j.physletb.2005.06.020),
[arXiv:hep-ph/0505013](https://arxiv.org/abs/hep-ph/0505013).
 - [95] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, “Phenomenology of GeV-scale heavy neutral leptons”, *JHEP* **11** (2018) 032,
[doi:10.1007/JHEP11\(2018\)032](https://doi.org/10.1007/JHEP11(2018)032), [arXiv:1805.08567](https://arxiv.org/abs/1805.08567).
 - [96] CMS Collaboration, “Measurement of the inclusive cross sections for W and Z boson production in proton-proton collisions at $\sqrt{s} = 5.02$ and 13 TeV”, [arXiv:2408.03744](https://arxiv.org/abs/2408.03744). (2024).

- [97] J. Beacham et al., “Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report”, *J. Phys. G* **47** (2020) 010501, [doi:10.1088/1361-6471/ab4cd2](#), [arXiv:1901.09966](#).
- [98] 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](#).
- [99] L. Canetti and M. Shaposhnikov, “Baryon Asymmetry of the Universe in the ν MSM”, *JCAP* **09** (2010) 001, [doi:10.1088/1475-7516/2010/09/001](#), [arXiv:1006.0133](#).
- [100] CHARM Collaboration, “A search for decays of heavy neutrinos in the mass range 0.5 - 2.8 GeV”, *Phys. Lett. B* **166** (1986) 473, [doi:10.1016/0370-2693\(86\)91601-1](#).
- [101] G. Bernardi et al., “Further limits on heavy neutrino couplings”, *Phys. Lett. B* **203** (1988) 332, [doi:10.1016/0370-2693\(88\)90563-1](#).
- [102] Belle Collaboration, “Search for heavy neutrinos at Belle”, *Phys. Rev. D* **87** (2013) 071102, [doi:10.1103/PhysRevD.87.071102](#), [arXiv:1301.1105](#). [Erratum: *Phys. Rev. D* **95** (2017) 099903].
- [103] LHCb Collaboration, “Search for Majorana neutrinos in $B^- \rightarrow \pi^+ \mu^- \mu^-$ decays”, *Phys. Rev. Lett.* **112** (2014) 131802, [doi:10.1103/PhysRevLett.112.131802](#), [arXiv:1401.5361](#).
- [104] 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](#).
- [105] 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”, *Phys. Rev. Lett.* **131** (2023) 061803, [doi:10.1103/PhysRevLett.131.061803](#), [arXiv:2204.11988](#).
- [106] 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) 221801, [doi:10.1103/PhysRevLett.120.221801](#), [arXiv:1802.02965](#).
- [107] 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](#).
- [108] SHiP Collaboration, “Sensitivity of the SHiP experiment to Heavy Neutral Leptons”, *JHEP* **04** (2019) 077, [doi:10.1007/JHEP04\(2019\)077](#), [arXiv:1811.00930](#).
- [109] FASER Collaboration, “FASER’s physics reach for long-lived particles”, *Phys. Rev. D* **99** (2019) 095011, [doi:10.1103/PhysRevD.99.095011](#), [arXiv:1811.12522](#).

-
- [110] 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](#).
 - [111] CMS Collaboration, “Search for heavy neutrinos and third-generation leptoquarks in hadronic states of two τ leptons and two jets in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **03** (2019) 170, [doi:10.1007/JHEP03\(2019\)170](#), [arXiv:1811.00806](#).
 - [112] CMS Collaboration, “Search for Long-Lived Heavy Neutral Leptons with Lepton Flavour Conserving or Violating Decays to a Jet and a Charged Lepton”, *JHEP* **03** (2024) 105, [doi:10.1007/JHEP03\(2024\)105](#), [arXiv:2312.07484](#).
 - [113] CMS Collaboration, “Search for long-lived heavy neutral leptons decaying in the CMS muon detectors in proton-proton collisions at $\sqrt{s}=13$ TeV”, *Phys. Rev. D* **110** (2024) 012004, [doi:10.1103/PhysRevD.110.012004](#), [arXiv:2402.18658](#).
 - [114] A.-M. Lyon, “Searches in CMS for long-lived particles and other non-conventional signatures”, *PoS (EPS-HEP2023)* (2024) 450, [doi:10.22323/1.449.0450](#).
 - [115] P. Baldi et al., “Parameterized neural networks for high-energy physics”, *Eur. Phys. J. C* **76** (2016) 235, [doi:10.1140/epjc/s10052-016-4099-4](#), [arXiv:1601.07913](#).
 - [116] A. Abada, P. Escribano, X. Marcano, and G. Piazza, “Collider searches for heavy neutral leptons: beyond simplified scenarios”, *Eur. Phys. J. C* **82** (2022) 1030, [doi:10.1140/epjc/s10052-022-11011-7](#), [arXiv:2208.13882](#).
 - [117] M. Drewes, J. Klarić, and J. López-Pavón, “New benchmark models for heavy neutral lepton searches”, *Eur. Phys. J. C* **82** (2022) 1176, [doi:10.1140/epjc/s10052-022-11100-7](#), [arXiv:2207.02742](#).
 - [118] A.-M. Lyon, “Review of searches for new physics at CMS”, in *58th Rencontres de Moriond on QCD and High Energy Interactions*. (2024). [arXiv:2406.02010](#).
 - [119] CMS Collaboration, “Search for long-lived heavy neutrinos in the decays of B mesons produced in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **06** (2024) 183, [doi:10.1007/JHEP06\(2024\)183](#), [arXiv:2403.04584](#).
 - [120] CMS Collaboration, “Review of searches for vector-like quarks, vector-like leptons, and heavy neutral leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV at the CMS experiment”, [arXiv:2405.17605](#). (2024), Accepted by *Phys. Rept.*
 - [121] LHCb Collaboration, “Measurement of the B_c^- meson production fraction and asymmetry in 7 and 13 TeV pp collisions”, *Phys. Rev. D* **100** (2019) 112006, [doi:10.1103/PhysRevD.100.112006](#), [arXiv:1910.13404](#).
 - [122] D. J. Lange, “The EvtGen particle decay simulation package”, *Nucl. Instrum. Meth. A* **462** (2001) 152, [doi:10.1016/S0168-9002\(01\)00089-4](#).

- [123] C.-H. Chang, J.-X. Wang, and X.-G. Wu, “BCVEGPY2.0: A Upgrade version of the generator BCVEGPY with an addendum about hadroproduction of the P-wave B(c) states”, *Comput. Phys. Commun.* **174** (2006) 241, [doi:10.1016/j.cpc.2005.09.008](https://doi.org/10.1016/j.cpc.2005.09.008), [arXiv:hep-ph/0504017](https://arxiv.org/abs/hep-ph/0504017).
- [124] CMS Collaboration, “Measurements of the HLT performance of displaced muons using Cosmics data”, (2019). <https://cds.cern.ch/record/2690807>.
- [125] K. Prokofiev and T. Speer, “A kinematic and a decay chain reconstruction library”, in *14th International Conference on Computing in High-Energy and Nuclear Physics*, p. 411. (2005).
- [126] CMS Collaboration, “CMS tracking performance results from early LHC operation”, *Eur. Phys. J. C* **70** (2010) 1165, [doi:10.1140/epjc/s10052-010-1491-3](https://doi.org/10.1140/epjc/s10052-010-1491-3), [arXiv:1007.1988](https://arxiv.org/abs/1007.1988).
- [127] M. Abadi et al., “TensorFlow: A system for large-scale machine learning”, in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI’16, p. 265. USENIX Association, USA, (2016). [arXiv:1605.08695](https://arxiv.org/abs/1605.08695).
- [128] F. Chollet et al., “Keras”. <https://keras.io>, (2015).
- [129] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization”, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980). (2014).
- [130] CMS Collaboration, “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018) P06015, [doi:10.1088/1748-0221/13/06/P06015](https://doi.org/10.1088/1748-0221/13/06/P06015), [arXiv:1804.04528](https://arxiv.org/abs/1804.04528).
- [131] CMS Collaboration, “Measurement of the total and differential inclusive B^+ hadron cross sections in pp collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **771** (2017) 435, [doi:10.1016/j.physletb.2017.05.074](https://doi.org/10.1016/j.physletb.2017.05.074), [arXiv:1609.00873](https://arxiv.org/abs/1609.00873).
- [132] P. D. Dauncey, M. Kenzie, N. Wardle, and G. J. Davies, “Handling uncertainties in background shapes: the discrete profiling method”, *JINST* **10** (2015) P04015, [doi:10.1088/1748-0221/10/04/P04015](https://doi.org/10.1088/1748-0221/10/04/P04015), [arXiv:1408.6865](https://arxiv.org/abs/1408.6865).
- [133] R. A. Fisher, “On the mathematical foundations of theoretical statistics”, *Phil. Trans. Roy. Soc. Lond. A* **222** (1922) 309, [doi:10.1098/rsta.1922.0009](https://doi.org/10.1098/rsta.1922.0009).
- [134] S. S. Wilks, “The large-sample distribution of the likelihood ratio for testing composite hypotheses”, *Annals Math. Statist.* **9** (1938) 60, [doi:10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360).
- [135] A. de Gouvêa and A. Kobach, “Global Constraints on a Heavy Neutrino”, *Phys. Rev. D* **93** (2016) 033005, [doi:10.1103/PhysRevD.93.033005](https://doi.org/10.1103/PhysRevD.93.033005), [arXiv:1511.00683](https://arxiv.org/abs/1511.00683).
- [136] B. Shuve and M. E. Peskin, “Revision of the LHCb limit on Majorana neutrinos”, *Phys. Rev. D* **94** (2016) 113007, [doi:10.1103/PhysRevD.94.113007](https://doi.org/10.1103/PhysRevD.94.113007), [arXiv:1607.04258](https://arxiv.org/abs/1607.04258).

- [137] R. Barouki, G. Marocco, and S. Sarkar, “Blast from the past II: Constraints on heavy neutral leptons from the BEBC WA66 beam dump experiment”, *SciPost Phys.* **13** (2022) 118, [doi:10.21468/SciPostPhys.13.5.118](https://doi.org/10.21468/SciPostPhys.13.5.118), [arXiv:2208.00416](https://arxiv.org/abs/2208.00416).
- [138] WA66 Collaboration, “Search for heavy neutrino decays in the BEBC beam dump experiment”, *Phys. Lett. B* **160** (1985) 207, [doi:10.1016/0370-2693\(85\)91493-5](https://doi.org/10.1016/0370-2693(85)91493-5).