A Search for new physics in events with a leptonically decaying Z boson and a large transverse momentum imbalance with the CMS Detector at the LHC

by

Nicholas Charles Smith

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy
(Physics)

at the

University of Wisconsin – Madison

2018

Defended on Monday August 27th, 2018

Dissertation approved by the following members of the Final Oral Committee:

Wesley H. Smith · Bjorn Wiik Professor of Physics

Sridhara R. Dasu · Professor of Physics

Matthew F. Herndon · Professor of Physics

Aki Hashimoto · Professor of Physics

Marshall F. Onellion · Professor of Physics

Abstract

A search for new physics in events with a leptonically decaying Z boson and a large transverse momentum imbalance in proton-proton collisions collected at $\sqrt{s} = 13\,\text{TeV}$ in 2016 with the Compact Muon Solenoid detector at the Large Hadron Collider is presented. The results of this search are interpreted in terms of a simplified model of dark matter production via spin-0 or spin-1 mediators, a scenario with a standard-model-like Higgs boson produced in association with the Z boson and decaying invisibly, a model of unparticle production, and a model with large extra spatial dimensions. No significant deviations from the background expectations are found, and limits are set on relevant model parameters, significantly extending the results previously achieved in this channel.

Acknowledgements

To my parents, John and Charlotte, thank you for your continuous support and encouragement over my lifetime, and for the excellent educational opportunities you made available to me, both in physics and elsewhere, without which I surely would not have found myself in this position. To my wife Lesley, thank you so much for your unwavering support, and for keeping me sane during the years we spent at CERN. To my advisor Wesley Smith, as well as to Sridhara Dasu and Matthew Herndon, thank you for the excellent guidance and wisdom you have given me, which tends to have ever-expanding applicability. To Alexander Savin, Maria Cepeda, and Bhawna Gomber, thank you for introducing me to the ecosystem of the CMS collaboration, and for your guidance in developing the analysis. To Tom Perry, Aaron Levine, Laura Dodd, Nate Woods, Devin Taylor, Tyler Ruggles, Kenneth Long, and Usama Hussain; thank you for the countless fruitful discussions throughout our time together. I could not have completed this work without the assistance of several thousand collaborators, and their enumeration on the author list of every publication by the CMS collaboration is exceptional evidence that scientific research, now more than ever, is a collaborative endeavor. In particular, to my collaborators on the mono-Z analysis team: Andreas Albert, Guillelmo Gomez-Ceballos, Dylan Hsu, Daniele Trocino, Ashraf Kasem, Chad Freer, and Darien Wood; thank you for your invaluable input and assistance in the completion of this work. To Aleš Svetek, Marcelo Vicente, and Pam Klabbers; thank you for your training and assistance in carrying out my hardware support duties, and for taking time to explain in detail how the various systems work, even when it was not obligatory. To all my friends in Madison, thank you so much for all the great memories from my first two years of graduate school.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Experimental high-energy particle physics is the study of the behavior of particles of matter at the shortest distance scales. This field has advanced over the last century to experimentally validate and iteratively refine the theoretical description of matter that forms a basis for the interpretation of results spanning distance scales from the cosmological to the subatomic. This description is known as the Standard Model (SM), a quantum field theory (QFT) that explains the behavior of three of the four known fundamental forces of nature, namely: the strong nuclear force, which is the dominant force responsible for the formation of atomic nuclei; electromagnetism, the force responsible for the formation of atoms, as well as the most visible force in our daily lives; and the weak nuclear force, which is the dominant force responsible for radioactive β decay of nuclei. The mathematical description of the strong nuclear force forms the quantum chromodynamics (QCD) sector of the SM, while the electromagnetic force and the weak nuclear force are described in unison by electroweak (EW) theory.

The principal experimental technique used in high-energy particle physics is the *scattering experiment*, where a beam of charged particles is accelerated (naturally

or artificially) to a high velocity and made to collide either with a fixed target or another particle beam. The typical result of this collision is that the particles will be deflected from each other by the electromagnetic force, or scatter *elastically*. An early example is Rutherford scattering, first carried out in the early 1900s [1], where a source of α particles (Helium nuclei, a species of radiation produced from nuclear decays) is pointed at a thin gold foil. The distribution of the deflection angle of the α particles was found to be consistent with a theory describing the atom as a very small central charge surrounded by a diffuse distribution of balancing charges, eventually identified as the nucleus and electron cloud. The distance scale probed by scattering experiments is limited by the kinetic energy of the beam: classically, as the electromagnetic field is conservative, the distance of closest approach r for a head-on collision of two particles is related to the beam energy as $E \propto 1/r$. The QFT description of particle interactions is more subtle, but the qualitative result is the same: to probe small distance scales, particle beams must be made to collide at high energy. Thus the campaign to produce ever more powerful accelerators was launched. Soon, the energy scale became sufficient to cause *inelastic* scattering to occur, where the incident particle momentum, and eventually even the particle number, was not conserved—matter-antimatter pairs were being produced through quanta of the fields of the fundamental forces, known as gauge bosons. The plethora of species of outgoing particles observed by these experiments was cataloged, and the patterns found in them gave rise to various conservation laws, eventually identified as symmetries of the SM. The behavior of the gauge bosons were probed by analyzing the rates of inelastic scattering as a function of the center of mass energy of the collision and of the momenta of the outgoing particles. From these data, a unified picture of subatomic matter was formed. Fundamental matter particles (fermions), such as electrons, positrons, and muons, interact with each other through gauge boson

mediators, such as photons and Z bosons, and acquire a mass through interactions with the Higgs boson. Composite particles, such as protons and neutrons, arise from self-interacting clumps of quarks and gluons. The particle content and interaction relations of the SM are described in detail in Section 2.1.

The SM has several deficiencies. The goal of the analysis presented here is to test theories of new physics beyond the SM that may resolve some of these deficiencies. Cosmological evidence in galactic rotation curves, galaxy cluster velocity distributions, and cosmic microwave background anisotropy indicates (Section 2.4) that approximately 25% of our universe is composed of some form of dark matter (i.e. matter that interacts very weakly with the SM). Some hypothesized extensions of the SM that include this dark matter are testable with particle colliders. The theoretical basis of the remaining fundamental force, gravity, despite having withstood experimental tests at distance scales from that of daily life to the size of the universe, remains separated from that of the SM. A theory of large extra dimensions (Section 2.5) may be able to reconcile the 17 order of magnitude discrepancy in the characteristic energy scale of general relativity and the electroweak scale, a necessary step in the unification of gravity and the SM. An exotic form of matter that cannot be described in terms of particles (Section 2.5), if observed, may provide important insight for our understanding of quantum field theories. The recently discovered Higgs boson is a critical component of electroweak theory, and its behavior has yet to be fully characterized (Section 2.6). In particular, the Higgs boson is not a stable particle, and the possibility for it to decay to new species of particle not experimentally detectable is predicted by some models of new physics.

The four most powerful particle accelerator/colliders constructed thus far are: the Large Electron-Positron Collider (LEP), which was in operation from 1989 until 2000, and collided beams of electrons and positrons at a center of mass energy up

to 209 GeV; the Hadron-Electron Ring Accelerator (HERA), which was in operation from 1992 until 2007, and collided protons with electrons/positrons at a center of mass energy of 318 GeV; the Tevatron, which operated from 1987 until 2011, and collided protons with antiprotons at a center of mass energy up to 1.96 TeV; and the Large Hadron Collider, presently in operation, which collides two proton beams at a center of mass energy up to 13 TeV. Each of these machines has provided significant input into the definition and validation of the SM. The LHC is now in its second period of operation, after a two-year shutdown to carry out upgrades. Run 1 of the LHC culminated in the discovery of the Higgs boson in 2012 [2, 3], predicted to exist almost fifty years prior. This analysis uses data collected in 2016—the first year of run 2 of the LHC—with the Compact Muon Solenoid (CMS) detector.

This analysis tests for the presence of new physics that is predicted to lead to the production of matter that is not directly detectable by CMS. Evidence for this production can be found through a net imbalance in the momentum of visible particles emanating from the collision point transverse to the collision axis, where the transverse momentum of the invisible matter is balanced by that of an unstable Z boson that subsequently decays to a pair of muons[1] or electrons that are detectable by CMS. This signature, referred to as the $\mathrm{Z}+p_{\mathrm{T}}^{\mathrm{miss}}$ or "mono-Z" final state, is well-suited to probe these models of new physics as it has relatively small and well-understood backgrounds. Prior experimental results, both at the LHC and at earlier colliders, investigated this same final state. However, this result obtains the best sensitivity thus far to the new physics models probed here. No significant excess over the background expectation is observed, and more stringent limits on the allowed parameter space of these models are found. The structure of this work follows closely that of Ref. [4], which is based on the work presented here.

[1] Muons are unstable particles, however their 2 μs lifetime makes them stable on the timescale of transiting the CMS detector.

Chapter 2

Theoretical models

2.1 The Standard Model

The Standard Model (SM) of particle physics is a quantum field theory (QFT) which can be described in its most concise form in terms of the components of its Lagrangian density:

$$\begin{aligned}\mathscr{L} = &-\frac{1}{4}F_{\mu\nu}F^{\mu\nu} \\ &+ i\bar{\psi}\not{D}\psi \\ &+ \bar{\psi}_i y_{ij}\psi_j\phi + \text{h.c.} \\ &+ |D_\mu\phi|^2 - V(\phi).\end{aligned} \tag{2.1}$$

There is a significant amount of implied information in this equation. Unpacking the equation fully and extracting the description of its components and how observable phenomena emerge from it can and does fill up entire textbooks (e.g. [5, 6, 7, 8]), and therefore we will only discuss a rather limited scope of the SM that is relevant to the discussion that follows in this text. It is still useful to step through each term of Eqn. 2.1 as a means of introducing the particle content of the SM. This equation

describes the SM configuration before electroweak symmetry breaking (EWSB), which we will discuss after introducing the last term.

The first term in Eqn. 2.1 represents the kinetic energy of the gauge fields. There are three families of gauge bosons that arise from quantization of these fields in the SM, corresponding to the generators of three symmetry groups under which Eqn. 2.1 is invariant.[1] They are the 8 gluon fields G of the $SU(3)$ Lie symmetry group that embodies the strong nuclear force, the 3 W fields of the $SU(2)_L$ group that embodies the weak isospin force, and the B field of the $U(1)$ weak hypercharge force. Each gauge group has an associated coupling strength g that is a free parameter of the SM, and appears in the self-coupling term encoded in this term if the group is non-abelian, otherwise it appears in the fermion coupling terms. The latter two symmetries become almost exact at energies significantly above the electroweak unification scale, however as we will see momentarily, the $SU(2)_L \times U(1)$ symmetry is spontaneously broken, resulting in a different symmetry and additional structure at low energies.

The second term represents the kinetic energy of the fermion fields ψ, as well as their coupling to the gauge fields. There are two classes of fermion fields, quarks q and leptons ℓ, where quarks are charged under the $SU(3)$ gauge group, and leptons are not. The fields can also be split based on their chirality: as the $SU(2)_L$ gauge group couples only to left-handed fermion fields, the left-handed fields are organized into doublets of up-type and down-type quarks or leptons, where the right-handed fields are uncharged under $SU(2)_L$. There are corresponding antifermions for all of the above classes, where the gauge group charge is inverted. To maintain the invariance under gauge transformations for the fields which transform non-trivially under a given gauge group, the gauge-covariant derivative $\not{D}$ is introduced, which causes interaction terms between the fermion fields and the gauge bosons to appear

[1] There are several additional symmetries that the SM Lagrangian respects, e.g. the Poincaré group that preserves momentum and Lorentz invariance.

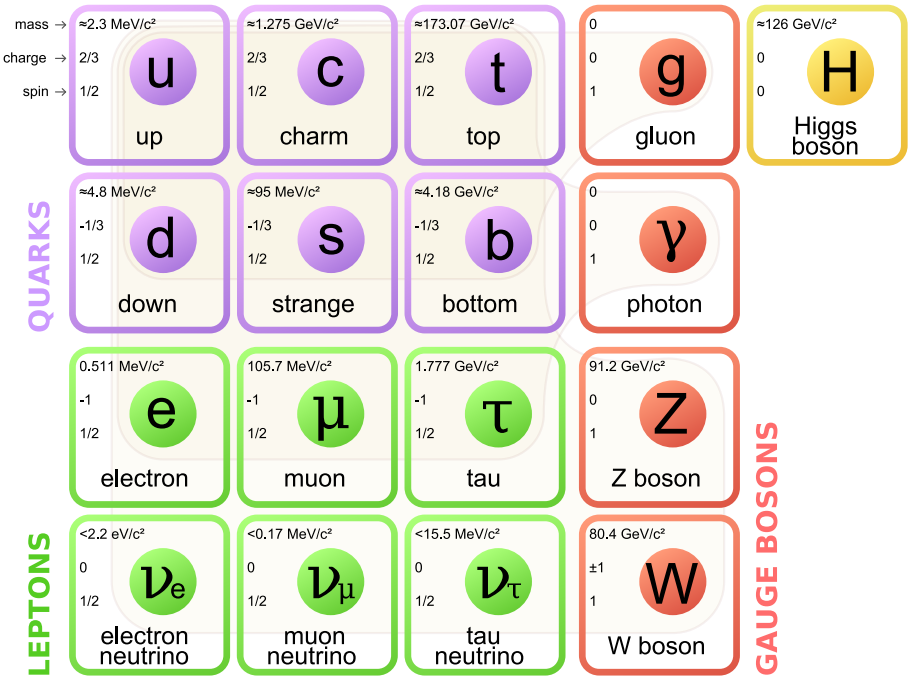

Figure 2.1: All fundamental particles in the Standard Model.

in the Lagrangian. This interaction generated by the symmetry is what causes forces to be communicated between matter, hence why the gauge bosons are referred to by the "force" they carry. Lastly, there are three identical copies of this arrangement, known as the three generations, or flavors, of matter. The fermion content of the SM is summarized in Fig. 2.1, albeit labeled with the post-EWSB gauge charges. The three generations of matter have identical properties except for their Yukawa couplings, which leads us to the next term.

The third term represents the coupling of the fermion fields to the complex scalar Higgs field ϕ. This type of term is known as a Yukawa term, originating from the Yukawa theory of nucleon binding, in which a pseudoscalar mediator (now identified as a composite particle, the pion) interacts with the nucleon with a similar Lagrangian description. The Yukawa coupling matrix is 3x3, running over the 3 generations of

matter, and there is a separate coupling matrix for quarks and leptons, and also a separate matrix for up-type and down-type fermions. These matrices represent the vast majority of free parameters in the SM, although some of the terms in the matrices are redundant or can be absorbed into redefinitions of the fields. The fourth term (hermitian conjugate) is a shorthand indicating that the same coupling terms are to be constructed also for the anti-fermions.

The fifth term represents the kinetic energy of the Higgs field, and, since ϕ is a doublet charged under $SU(2)_L$, also its interaction with the W fields. The sixth term $V(\phi)$ is the potential energy of the Higgs field. The full potential is

$$V(\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2$$

where μ and λ are free parameters of the SM. The Higgs potential energy is one of the most crucial pieces of the SM Lagrangian, since if $\mu^2 < 0$, a mechanism known as electroweak symmetry breaking occurs: $V(\phi)$ will have a minimum when $\phi^\dagger\phi = -\mu^2/2\lambda$, hence one of the four components of ϕ has a nonzero vacuum expectation value $v = \sqrt{-\mu^2/\lambda}$. If the real component with weak isospin $-1/2$ is chosen, then since a non-observable $SU(2)$ rotation can fix the remaining components to zero, we are left with one degree of freedom, which is identified as the Higgs boson H. By expanding any term in the SM Lagrangian involving ϕ into two terms with

$$\phi = \begin{pmatrix}\phi^+\\ \phi^0\end{pmatrix} \rightarrow \begin{pmatrix}0\\ v\end{pmatrix} + \begin{pmatrix}0\\ h\end{pmatrix},$$

the structure of the Lagrangian is significantly altered. In particular, the W and B fields are rotated into a new basis in which three massive bosons, the $\mathrm{W}^\pm$ and Z bosons, and one massless boson, the photon γ, appear. The photon becomes the gauge boson of a remnant $U(1)$ symmetry to be identified with electromagnetism, under which the $\mathrm{W}^\pm$ bosons as well as all of the fermions except neutrinos are charged. In addition, the Yukawa couplings of the fermion fields become fermion mass terms,

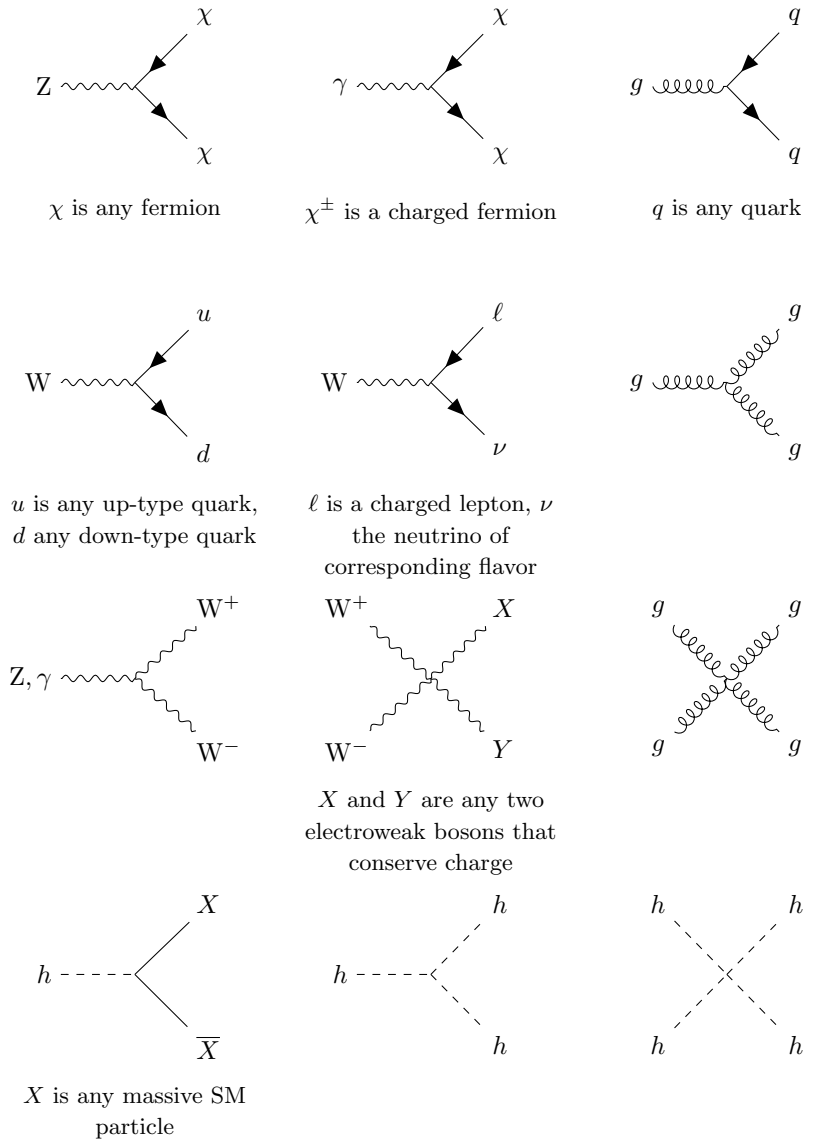

Figure 2.2: All interactions allowed between particles in the Standard Model.

along with couplings to the Higgs boson. Thus, the SM particle content obtains the properties shown in Fig. 2.1, and the interactions shown in Fig. 2.2.

The diagrams shown in Fig. 2.2 are more than schematic: the quantum field theory framework on which the SM is built allows the conversion from terms (also referred to as operators) in the Lagrangian to Feynman rules (formulas) associated with each

vertex and line in the diagrams. These rules/diagrams can be pieced together to build amplitudes for scattering processes, which upon summing and squaring, gives real-valued probabilities. Several diagrams can be constructed with identical incoming and outgoing particles, each with successively higher numbers of vertices, where each vertex contributes a weight proportional to the coupling g associated with the gauge interaction that gives rise to the vertex. The diagram for a given process with the lowest possible multiplicity of vertices is known as the *leading-order* diagram. The diagram with the next lowest possible multiplicity is called the *next-to-leading-order* (NLO) diagram, and so forth. In reality, all possible diagrams will contribute to the scattering probability, however one typically cannot evaluate this infinite sum of diagrams. For small g, the weight of diagrams with additional vertices is suppressed compared to that of the LO diagram, and in this sense the higher order diagrams can be treated as small corrections to the scattering amplitude of the LO diagram. This treatment is known as perturbation theory, and a necessary assumption is that the series of higher order corrections will converge. This is true if g is sufficiently small, or *perturbative*. The demarcation point at which a coupling g is considered perturbative is not exact, as the combinatoric multiplicity of the diagrams needs to also be considered, but generally $g \lesssim 4\pi$ is considered perturbative.

One side-effect of this theory is that one can construct diagrams with closed loops (i.e. a cycle in the graph of the diagram) which have an associated momentum vector that is unconstrained, and can be integrated over all momentum space, often resulting in infinite scattering probabilities. These are obviously unphysical, and a scheme to remove these infinities by balancing them with new terms in the Lagrangian, known as renormalization, is introduced. Introducing renormalization creates alterations to mass and coupling terms in the Lagrangian, changing them from "bare" to renormalized parameters, and also creates a dependence on some renormalization energy

scale. This scale is unobservable, however it does describe how effective magnitudes of these parameters evolve with energy. The equation governing this evolution is known as the renormalization group equation. In the case of gauge couplings, this equation is dependent on the gauge group and the field content of the theory. In particular, the $SU(3)$ coupling of the strong force g_s effectively becomes larger at lower energy scales, such that eventually the particles interacting via the strong force become bound together, which is known as confinement.

The confinement property of the strong force is what gives rise to a zoo of composite subatomic particles known as hadrons. In fact, free quarks cannot be observed at all, and are bound together by gluons according to the requirement that the $SU(3)$ charge (known as color charge) is zero in the composite object, so that it decouples from the gluon field when in its final state.[2] Mesons are combinations of quarks $q\bar{q}$, where the color charge is canceled by the equal and opposite charge of the antiquark. Baryon are combinations of quarks qqq where each has a different color charge, and the structure of the $SU(3)$ algebra gives a net zero color charge. The binding energy of these states is often significant, giving hadrons a mass typically much larger than that of the constituent quarks. Heavier hadrons can decay into lighter hadrons, or to leptons and photons if a coupling exists and is kinematically allowed. Hadrons of particular importance include: protons $p = uud$, neutrons $n = udd$, charged pions $\pi^+ = u\bar{d}$ and $\pi^- = \bar{u}d$, and neutral pions $\pi^0 = (u\bar{u} - d\bar{d})/\sqrt{2}$. The characteristic scale of confinement corresponds to the scale at which g_s is no longer perturbative, about 1 GeV. The mass spectrum of hadrons dies off not far above this scale, around the 10 GeV range.

Another emergent property of color confinement is the evolution from free partons

[2] The strong force still affects hadrons at low energy, as it is responsible for binding of nuclei in atoms. However, it is communicated not by the gluon field but by an emergent pion field, which encodes a low energy approximate $SU(2)$ symmetry known as isospin.

(quarks or gluons) created at a high energy scale to hadrons which are stable enough to be detected. This process is known as *hadronization*, and will be discussed further in Section 5.5. An important feature of hadronization is that the kinematic distribution of produced hadrons is concentrated in the directions of the initial partons. These localized concentrations of hadrons are known as *jets*.

Thus we have built a 19-parameter[3] model of the subatomic particle content of the SM and their interactions. A likely question is whether or not this model has anything at all to do with reality. The answer, as understood so far, appears to be that it almost entirely explains the experimental observations collected over the last century, starting perhaps from the Rutherford scattering experiments that exposed the nucleus of the atom [1], up to the observation of the last fundamental particle of Fig. 2.1, the Higgs boson, in 2012 [2, 3]. A thorough review of present state of particle physics is published biannually by the Particle Data Group as the *Review of Particle Physics*, the most recent of which is [9]. Deficiencies in the SM, both mathematical and physical, do exist. Rather that attempt an exhaustive list of such deficiencies, we will discuss those which provide motivations for the Beyond the Standard Model (BSM) theories probed in this analysis in their respective sections.

2.2 The initial and final states

In this analysis, we explore the possibility of extensions to the SM that may produce a particular signature in the CMS detector at the LHC. What this means from a theoretical perspective, is that we have a series of events with an initial state and a possible final state, and we would like to know with what probability the interaction will occur. In these events, our initial state is two protons colliding head-on at a

[3]Or 26, if considering neutrino masses.

center-of-mass energy $\sqrt{s} = 13\,\text{TeV}$, and in the vast majority of cases, the protons pass by each other without interacting. In some cases, the protons *scatter elastically*, where the composite nature of the proton is maintained. Sometimes, however, an *inelastic scattering* occurs, where a fundamental parton from a proton interacts directly with a corresponding parton from the other proton, and produces one of any number of possible final states. Here, we are interested in a particular final state, where two opposite-sign same-flavor light leptons (i.e. e^+e^- or $\mu^+\mu^-$), which have an invariant mass compatible with the Z boson mass, are found in the detector along with a substantial momentum imbalance in the plane transverse to the proton-proton collision (hereafter, $p_{\text{T}}^{\text{miss}}$.) (N.B. the initial proton-proton state has approximately no transverse momentum in this frame of reference.) For a given final state, the probability of observing an event will be proportional to what is known as the cross section for this process, σ. The expected number of events observed in the CMS detector over a given time frame is given by $\sigma \int \mathcal{L} dt$, where $\mathcal{L}(t)$ is known as the instantaneous luminosity, which is a function of the operating characteristics of the LHC, and will be discussed further in Section 4.2.

A cross section is ultimately an integration of a matrix element $|\mathcal{M}|$, consisting of summed amplitudes from the allowed Feynman diagrams of a given process, as a function of the incoming and outgoing particle states $\mathbf{k}$, over the Lorentz-invariant 4-momentum phase space $d\Phi_n$ of the n incoming and outgoing objects [9]:

$$\sigma = \int |\mathcal{M}(\mathbf{k}_1, \ldots, \mathbf{k}_n)|^2 d\Phi_n(\mathbf{k}_1, \ldots, \mathbf{k}_n) \tag{2.2}$$

In the case of inelastic proton collisions, since the proton is not a fundamental particle but rather made up of partons, some fraction x of the total proton energy is carried by each parton that interacts, hence the energy of the actual interaction is given by

$$\sqrt{\hat{s}} = x_1 x_2 \sqrt{s}. \tag{2.3}$$

To properly calculate the total cross section, we need a method to sum over the possible incoming parton states as well as integrate over the distribution of their momentum fractions in the proton. This is accomplished using parton distribution functions (PDFs), which will be discussed further in Section 5.2.

Often, we are also interested in the *differential cross section*, $d\sigma/dp$ for some kinematic parameter of interest p. In the case of event simulation, the differentiation is total in the outgoing state kinematics, and simulated events are created by sampling the probability distribution described by $|\mathcal{M}|^2$. More details pertaining to the simulation of events are presented in Chapter 5.

2.3 SM backgrounds

There are many SM processes which create a final state identical to that of our hypothetical signal. Thanks to decades of experimental and theoretical work, the SM has an impressive and well-validated predictive power, which allows us to know—with some, ideally small, uncertainty—the amount of background events we will find in an experiment, as will be discussed extensively in Chapter 7. For now, it suffices to know that all SM backgrounds can be divided into two categories: *reducible background*, where the final state differs from that of the signal process, but may end up appearing identical due to deficiencies in the experimental apparatus; and *irreducible background*, where the final state is identical to that of the signal process, and the only method of discrimination comes from our *a priori* knowledge of the SM expectation, both in terms of total cross section and the differential cross section with respect to a parameter of interest. Often, the differential cross section of the irreducible SM background will differ from that of the signal process, which allows better discrimination between the two in the experimental analysis. In the $\mathrm{Z} + p_{\mathrm{T}}^{\mathrm{miss}}$ final

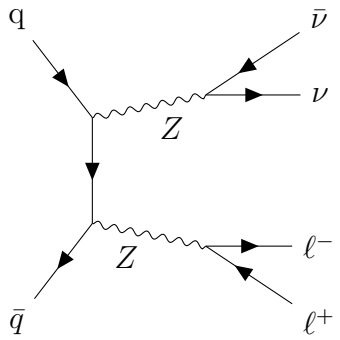

Figure 2.3: Feynman diagram showing the leading-order production mechanism for the ZZ $\rightarrow 2\ell 2\nu$ SM background process. This background is irreducible, as it has the same final state configuration as the signal models explored in this analysis, namely two opposite sign same flavor leptons in addition to a significant transverse momentum imbalance originating from the undetectable neutrinos.

state, there is only one irreducible background, the ZZ $\rightarrow 2\ell 2\nu$ process, for which the leading-order diagram is shown in Fig. 2.3. The calculation of the expectation of this background, as well as that of several other backgrounds, will be discussed further in Section 5.4. The reducible backgrounds for this final state will be discussed in more detail in Section 7.2.

2.4 Simplified particle dark matter models

2.4.1 Motivation for simplified particle dark matter

The nature of dark matter (DM) is one of the most significant puzzles in modern physics. Evidence has compounded over the last century for the existence of a type of matter that interacts with the matter content of the SM only very weakly, yet has mass and hence gravitates. The extreme weakness of gravity restricts currently known observed effects of DM to only the largest mass and distance scales, and evidence for DM can be found at distance scales starting from that of a galaxy.

Evidence for dark matter at the galactic level is found in the observation [10, 11]

that the rotational velocity of stars in galaxies does not behave as expected based on the visible mass distribution. From an equilibrium two-body treatment of the star-galaxy system, stars can be expected to orbit the galaxy center with a velocity $v(r) = \sqrt{GM(r)/r}$, where r is the radial distance from the center of the galaxy, G is the gravitational constant, and $M(r)$ is the total mass enclosed by the orbit. Stars far from the center of galaxies, where the visible mass density is much lower than the central mass density, would perceive an approximately constant $M(r)$ and thus $v \propto 1/\sqrt{r}$. Rather, it is observed (see Fig. 2.4) that $v(r)$ becomes constant in r towards the edges of galaxies, suggesting that significant additional non-visible mass is present. Analyses including more sophisticated treatments of the dynamics using multi-body simulations reach similar conclusions [12].

Additional evidence for dark matter can be found at the galactic cluster level. In particular, it was shown [13] in 1933 that the observed relative velocities of galaxies in the Coma cluster suggest a mass to light ratio over 2 orders of magnitude higher than expected. Modern analyses include also observations of gravitational lensing by clusters and X-ray emission from hot gas in dense clusters, and can provide direct measurements of the dark matter density at the ten percent level [12]. It is worth noting that these direct observations of galaxy and galactic cluster motion may also be explained by deviations from the general relativity framework of gravity at large distance scales. Such modified gravity theories would obviate the need for dark matter. However, the remarkable success of general relativity from distance scales of 1 mm to the size of the solar system makes acceptance of modified gravity difficult.

Although observations at galaxy and galactic cluster scales provide ample evidence for the existence of dark matter, the best quantitative prediction of the relative abundance of dark matter in the universe comes from global fits of several data sources at the cosmological scale to the Λ_{CDM} model, the so-called "Standard Model of Cos-

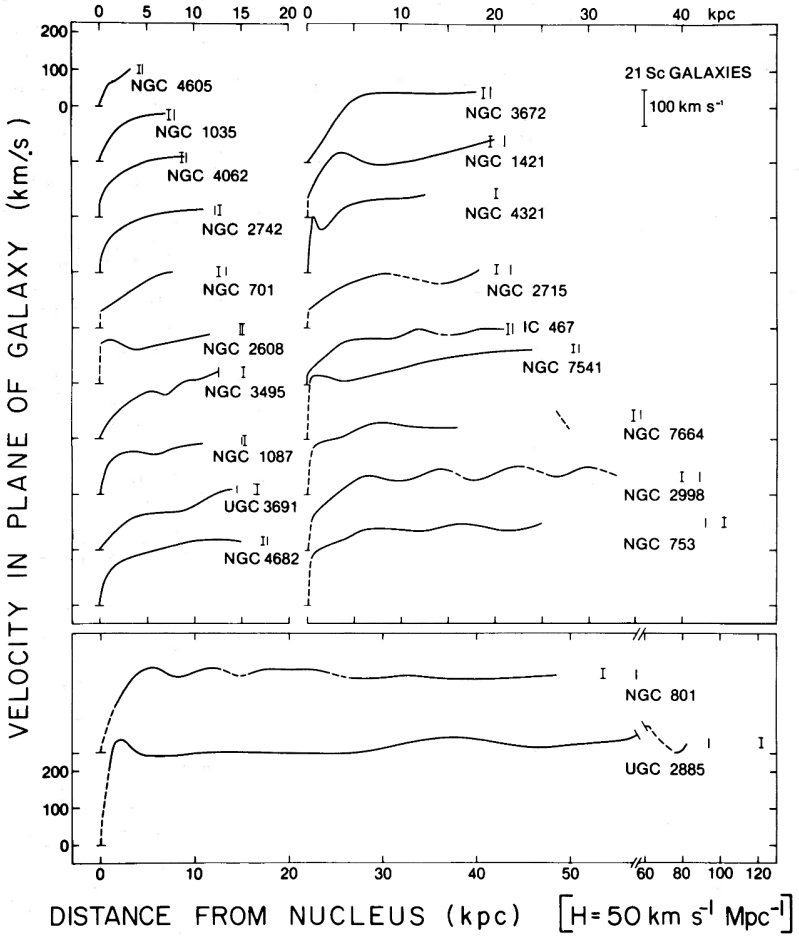

Figure 2.4: Mean velocities in the plane of the galaxy, as a function of linear distance from the nucleus for 21 Sc-class (spiral galaxies with loosely-wrapped arms) galaxies, arranged according to increasing linear radius. Reproduced from [10], Fig. 5.

mology" ([9], Chapter 22 contains a nice review.) The development of the Λ_{CDM} model could arguably have started as soon as the framework of general relativity was established in 1915 [14]. A solution to the Einstein field equations, known as the Friedmann–Lemaître–Robertson–Walker (FLRW) metric [15], describes a homogeneous isotropic expanding[4] universe, which—although at first glance appears to be a terrible approximation—is in fact a highly accurate description of our universe at

[4] A contracting universe is also possible, and the sign of the curvature parameter k in the Λ_{CDM} model controls whether the universe is open or closed. Evidence points towards $k = 0$, corresponding to a flat universe (which will still expand.)

cosmological distance scales: the cosmic microwave background (CMB) is measured to be isotropic within one part in 10^5 [16], and galaxy surveys show a homogeneous distribution at distance scales above 10^8 light years [17]. The FLRW metric leads to a set of two differential equations, known as the Friedmann equations, which when coupled with a thermodynamic model of the matter and energy density and pressure in the universe, allows quantitative predictions about the time evolution of the geometry of the universe and its contents to be made. Each component of the universe is represented by a time-dependent density parameter Ω, and a pressure-to-density parameter that governs how it evolves in time. These components interact, although the thermodynamic nature of the model restricts interactions to when the temperature of the universe is above the relevant energy scale of the interaction. As the universe expands, the temperature decreases, successively "freezing-out" components as they become decoupled from thermal equilibrium. These decoupled components are also referred to as thermal relics. The $\Lambda_{\rm CDM}$ model divides content into: dark energy, or equivalently, a cosmological constant Ω_Λ; cold (i.e. non-relativistic) dark matter Ω_c; baryonic (atoms) matter Ω_b; radiation (e.g. the CMB) Ω_r; and relativistic matter such as neutrinos, Ω_ν. With these pieces, one can utilize the $\Lambda_{\rm CDM}$ model to tie measurements of the CMB anisotropy, galactic surveys, supernovae distance measurements, and other cosmological observations together in one framework to find, among other parameters, the present-day relative abundance of visible matter, dark matter, and dark energy in the universe to be $\Omega_b = 0.0484(10)$, $\Omega_c = 0.258(11)$, and $\Omega_\Lambda = 0.692(12)$, respectively [18, 19].

Due to the cosmological evidence for DM outlined above, an entire experimental industry has built up focused around finding additional direct evidence for DM. The $\Lambda_{\rm CDM}$ model asserts that DM was in thermal equilibrium with baryonic matter at some point in the history of the universe, therefore it must interact with SM

particles at some energy scale. Several candidate particles have been proposed, the most popular of which is the weakly interacting massive particle, or WIMP. The WIMP interaction with the SM is assumed to occur at or around the weak force scale ($\sim$0.1-1 TeV), and the WIMP mass is assumed to be in the GeV range. The weak interaction assumption is motivated by the observation of the thermal relic density: once dark matter becomes decoupled from thermal equilibrium, the time evolution of Ω_c becomes fixed by the Friedmann equations, and therefore measurements of the present abundance can inform the temperature (energy scale) at which it decoupled, setting the energy scale at which interaction with baryonic matter can be expected. The assumed mass range of the WIMP is motivated by the necessity that Ω_c is sufficiently non-relativistic (cold) after decoupling such that it has the correct equation of state, as relativistic particles have a different pressure to density ratio. DM masses much lower than the GeV range can be accommodated in the Λ_{CDM} model, albeit with some modifications to the interaction model [20]. In the WIMP dark matter paradigm, detection techniques can be broadly classified into the three mechanisms (Fig. 2.5) through which DM can interact with the SM: indirect detection of DM annihilation via resonances in cosmic ray $(p, e^{\pm}, \gamma, \ldots)$ energy distributions; direct detection of DM scattering against atomic nuclei; and pair production from SM particle annihilation at colliders. A thorough review of both direct and indirect detection experimental results is available in [12]. A review of previous collider results will be presented in Section 3.2, along with relevant direct and indirect detection results.

2.4.2 Simplified particle DM model description

The framework in which collider limits on DM production were first explored was through effective field theory [21, 22, 23, 24]. The concept of effective field theory (EFT) is to reduce the model complexity of a full theory with a high characteristic

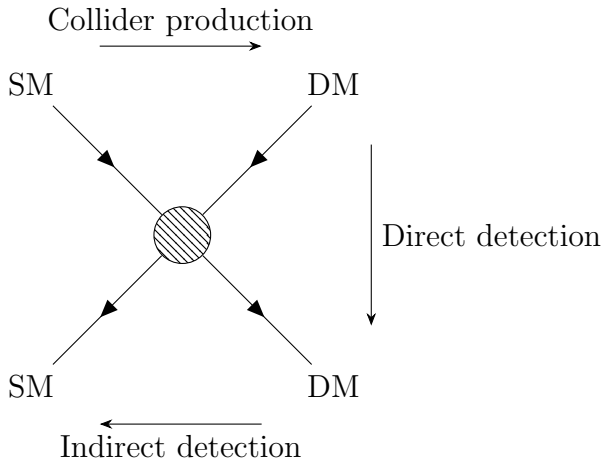

Figure 2.5: Schematic diagram of particle dark matter detection techniques.

energy scale Λ by making certain approximations that hold if the energy scale of the interaction $\sqrt{\hat{s}} \ll \Lambda$. A classic example is Fermi theory [25], which describes beta decay of neutrons via a single contact operator between a nucleon, electron, and neutrino field with a coupling parameter G_F that we now know to be [6]

$$G_F = \frac{\sqrt{2}g^2}{8m_{\mathrm{W}}^2},$$

where g is the $SU(2)_L$ weak coupling constant and m_{W} is the mass of the W boson. This relation can be derived from the resolved process in the full SM,

$$udd(n) \to uud(p) + \mathrm{W}^- \to uud(p) + e^- + \bar{\nu}_e,$$

upon integrating out the W boson propagator. Fermi theory is valid for this process because the energy scale of the interaction, $m_n - m_p = 1.29\,\mathrm{MeV}$, is much lower than the scale of the resolved process, $m_{\mathrm{W}} = 80.4\,\mathrm{GeV}$ [9]. The power of EFT comes in the ability to make meaningful predictions, as Fermi did, without knowledge of the high-energy theory. As a bonus, the cutoff scale gives us some idea at which energies we will see new physics. Indeed, the W boson was predicted [26] before it was found [27, 28].

In the EFT treatment of collider DM production, for a given hypothesized dark matter species (typically assumed to be fermionic χ), a full list of operators coupling DM to the SM is produced (see e.g. Table 1 of [24]) from which the production phenomonology (in particular, differential distributions) of terms with dominant amplitudes is investigated. In many early LHC DM analyses [29, 30], the differential distributions of these EFT operators were used to set limits on the EFT cutoff scale which were typically only 1-10 times the center-of-mass energy of the collision, at which point one would naturally expect signs of new physics to emerge directly, e.g. in high-mass resonance searches. This apparent paradox has been noticed [31, 32], and has prompted the building of models that contain a resolved mediator between DM and the SM [33, 34], trading model-independence for more meaningful results. Four such models are explored here, where a fermionic dark matter candidate with mass m_{DM} couples to a mediator of mass m_{med} with coupling g_{DM}, which in turn couples to SM quarks with coupling g_q. The mediator can take one of four spin-parity representations: vector or axial vector, denoted Z'; and scalar or pseudoscalar, denoted ϕ. The respective operators in the DM Lagrangian are:

$$\begin{aligned}
\mathscr{L}_{\text{vector}} &= g_{\mathrm{DM}} {Z'}_{\mu} \bar{\chi} \gamma^{\mu} \chi + g_q \sum_q {Z'}_{\mu} \bar{q} \gamma^{\mu} q, \\
\mathscr{L}_{\text{axial vector}} &= g_{\mathrm{DM}} {Z'}_{\mu} \bar{\chi} \gamma^5 \gamma^{\mu} \chi + g_q \sum_q {Z'}_{\mu} \bar{q} \gamma^5 \gamma^{\mu} q, \\
\mathscr{L}_{\text{scalar}} &= g_{\mathrm{DM}} \phi \bar{\chi} \chi + g_q \frac{\phi}{\sqrt{2}} \sum_q y_q \bar{q} q, \\
\mathscr{L}_{\text{pseudoscalar}} &= i g_{\mathrm{DM}} \phi \bar{\chi} \gamma^5 \chi + g_q \frac{i\phi}{\sqrt{2}} \sum_q y_q \bar{q} \gamma^5 q.
\end{aligned} \tag{2.4}$$

These models are to be treated as benchmark theories, and do not provide any information about new physics at energy scales above m_{med}. In particular, the scalar mediator will mix with the SM Higgs field, causing alterations of Higgs sector physics that may be difficult to reconcile with present experimental results; and, although

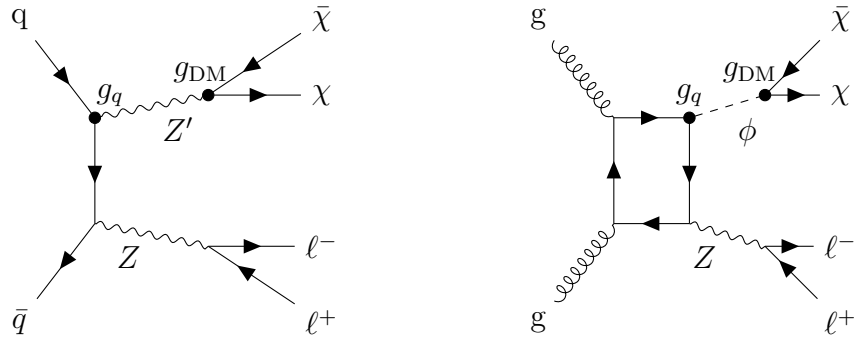

Figure 2.6: Feynman diagrams showing leading-order production of dark matter particles in association with a Z boson: (left) vector or axial vector mediated production; (right) scalar or pseudoscalar mediated production.

renormalizable, the vector models require a mechanism[5] for the vector mediator to acquire mass, and additional dark sector fermion fields to cancel the anomaly [34].

To detect pair production of DM in these models at the LHC, an additional SM interaction is required: if the result of the collision is two invisible particles and nothing else, there is no way to detect the process. The most common form of additional interaction is initial state radiation (ISR), where a SM boson (typically a gluon, photon, or Z boson) is radiated from one of the incoming quarks. The leading-order diagrams showing the production of DM in association with an ISR Z boson are shown for these models in Fig. 2.6.

For the vector and axial vector mediated DM models, after fixing the coupling strengths g_q and g_{DM}, collider constraints on the $(m_{\mathrm{DM}}, m_{\mathrm{med}})$ parameter space can be translated into constraints on the DM-nucleon cross section probed by direct detection

[5] A dark Higgs sector, with a new Higgs doublet that only interacts with the DM mediator and/or the DM fermions, can provide a mass for the DM mediator in the same way as the SM Higgs provides mass to the Z boson.

experiments (see Section 3.2.) The exact relation is given [35]:

$$
\begin{aligned}
F(g_q, g_{\mathrm{DM}}, m_{\mathrm{med}}, m_{\mathrm{DM}}) &= \left(\frac{g_q g_{\mathrm{DM}}}{0.25}\right)^2 \left(\frac{1\,\mathrm{TeV}}{m_{\mathrm{med}}}\right)^4 \left(\frac{m_n m_{\mathrm{DM}}}{(m_n + m_{\mathrm{DM}}) \cdot 1\,\mathrm{GeV}}\right)^2, \\
\sigma_{\mathrm{SI}} &= 6.9 \times 10^{-41}\,\mathrm{cm}^2 \cdot F(\ldots), \quad \text{(vector)} \\
\sigma_{\mathrm{SD}} &= 2.4 \times 10^{-42}\,\mathrm{cm}^2 \cdot F(\ldots), \quad \text{(axial vector)}
\end{aligned}
\tag{2.5}
$$

where $m_n = 0.939\,\mathrm{GeV}$ is the average mass of a nucleon. Collider constraints on scalar and pseudoscalar mediated DM models can also be compared to direct and indirect detection experiments via similar relations. Another common re-interpretation included in presentations of the allowed $(m_{\mathrm{DM}}, m_{\mathrm{med}})$ parameter space for these models is the region which gives a DM thermal relic abundance $\Omega_c h^2 \leq 0.12$. The inequality represents the possible presence of an additional DM species in the universe, which may contribute additively to the DM energy density [35].

2.5 Large extra dimension and unparticle models

2.5.1 Motivation for the models

The Arkani-Hamed–Dimopoulos–Dvali (ADD) model with large extra spatial dimensions [36, 37, 38] is motivated by the hierarchy problem, namely, the disparity in the magnitudes of the electroweak unification scale $m_{\mathrm{EW}} \approx 1\,\mathrm{TeV}$ and the Planck scale $m_{\mathrm{Pl}} \approx 10^{16}\,\mathrm{TeV}$. In the ADD model, n compact extra dimensions exist which have a large characteristic size compared to the electroweak scale, in which Kaluza-Klein gravitons (G) can propagate. For energies below the electroweak scale, SM particles are confined to propagate in the normal 4-dimensional space. Then the apparent weakness of gravity (or, equivalently, the apparent largeness of m_{Pl}) is explained by the fact that gravitons can propagate in a higher-dimensional space.

The unparticle model [39] consists of a hypothetical field theory which becomes scale-invariant at low energy and interacts with the SM at high energy. Scale invariance is a type of symmetry that is of theoretical interest because any model with this property would automatically continue to be valid (from a computational standpoint, not necessarily the *correct* description of nature) to arbitrarily small length scales, or equivalently, arbitrarily high energies. Scale invariance is not present in the SM due to the only dimensionful parameter in the (pre-EWSB) Lagrangian: the μ parameter of the Higgs potential $V(\phi)$. The unparticle model is motivated not by any particular issue in the SM, but rather simply because it explores an unusual quantum field theory which does not have any concept of free particle states with a well-defined nonzero mass, which by itself can be of theoretical interest, but also makes testable predictions.

2.5.2 ADD and unparticle model description

In the ADD model, the apparent Planck scale is related to the true Planck scale M_{D} (to be considered synonymous with the electroweak scale in this theory) as [36]

$$m_{\mathrm{Pl}}^2 = \mathcal{C} M_{\mathrm{D}}^{2+n} R^n \tag{2.6}$$

where R is the size of the compact extra dimensions, and $\mathcal{C}$ is some geometric factor determined by the topology of the extra dimensions.[6] To recover the apparent Planck scale, R is fixed by this relation, and the free parameters in the theory are M_{D} and n. The ADD model is realized by forming the $(n+4)$-dimensional EFT Lagrangian for a Graviton and including a coupling to the SM energy-momentum tensor in [40], and is ultimately reduced to a matrix element for the process $\mathrm{q\bar{q}} \to Z + G$ (Fig. 2.7.) Since the graviton model is an effective field theory, it can only be valid below some

[6] In [36], the extra dimensions were assumed to form a n-dimensional torus, which gave a geometric factor of $(2\pi)^n$

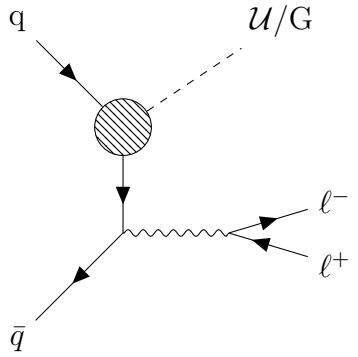

Figure 2.7: Feynman diagram illustrative of the production of Gravitons in the ADD model or unparticles.

cutoff scale. From dimensional arguments, the scale at which perturbative expansion breaks down is $\Lambda > 7.2M_{\mathrm{D}}$ ([40], Eqn. 68) and thus in this analysis a truncation procedure [41] is used whereby any portion of the signal phase space with $\sqrt{\hat{s}} > M_{\mathrm{D}}$ is suppressed by a weighting factor

$$w = \frac{M_{\mathrm{D}}^4}{\hat{s}^2} \tag{2.7}$$

In the description of the unparticle model, no assumptions are made about the structure of the theory other than scale-invariance at low energy, although the Banks–Zaks field [42] presents an example of a QCD-like theory where a particular particle structure can give a renormalization group equation for the coupling constant which has an infrared fixed point. The interaction term in the Lagrangian that produces the unparticles, in which the initial state is a pair of quarks, is sketched as

$$\mathscr{L}_{\mathcal{U}} = \frac{\lambda}{\Lambda_{\mathcal{U}}^{d_{\mathcal{U}}-1}} q\bar{q}O_{\mathcal{U}} \tag{2.8}$$

where λ describes the SM-unparticle coupling strength, Λ the EFT cutoff scale, $d_{\mathcal{U}}$ the scale dimension parameter, and $O_{\mathcal{U}}$ a spin-0 operator in the new theory. In this analysis, we consider only the case where the unparticle operator is spin-0, although alternative options have been explored [43]. The main phenomenological consequence

of the SM interaction with the unparticle sector is that the phase space factor is found to be the standard $d\Phi_n$ of Eqn. 2.2, except that $n = d_{\mathcal{U}}$ is a non-integer number of massless particles. Following the production of the unparticles, the weakness of the interaction is assumed to allow them to evade detection. If production is accompanied by an ISR Z boson, as depicted in Fig. 2.7, this creates the $\mathrm{Z} + p_{\mathrm{T}}^{\mathrm{miss}}$ final state.

2.6 Invisible Higgs boson decay models

2.6.1 Motivation for invisible Higgs decays

After the discovery of a Higgs boson [2, 3], one of the main experimental programs at the LHC is to measure the couplings of this boson to other SM particles, and compare these measurements to the SM predictions. Using all of the data collected in the first running period of the LHC, the CMS and ATLAS collaborations created a combined measurement of the Higgs couplings [44]. In this measurement, the SM predictions for both the production mechanisms and branching fractions were re-parameterized in terms of coupling modifiers κ, which were then constrained by the data. By including also a free parameter allowing some amount of the produced Higgs bosons to decay into some non-SM final state, B_{BSM}, the data shows that there is substantial room for BSM decays, with a 95% CL upper limit of $B_{\mathrm{BSM}} \leq 34\%$. The best fit value of B_{BSM}, as well as the other Higgs coupling modifiers, is shown in Fig. 2.8.

One possible source of B_{BSM} is invisible final states, i.e. decays to any particles that would not be detectable except by the transverse momentum imbalance they would cause. The possibility of the Higgs boson decaying to invisible final states (hereafter, H(inv.)) is directly measurable in collider experiments in multiple final states, one of which is the $\mathrm{Z} + p_{\mathrm{T}}^{\mathrm{miss}}$ final state explored here. The ZH(inv.) production mechanism is shown in Fig. 2.9.

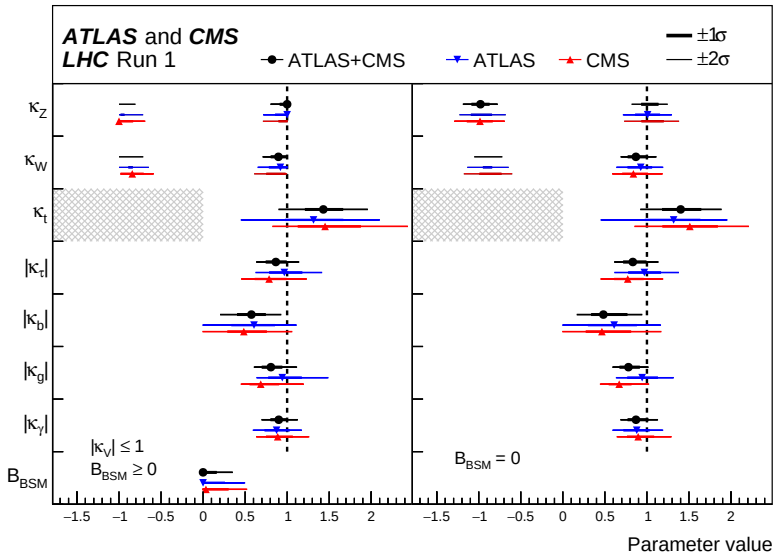

Figure 2.8: Fit results for two parameterizations allowing BSM loop couplings discussed in the text: the first one assumes that $B_{\mathrm{BSM}} \geq 0$ and that $|\kappa_{\mathrm{V}}| \leq 1$, where κ_{V} denotes κ_{Z} or κ_{W}, and the second one assumes that there are no additional BSM contributions to the Higgs boson width, i.e. $B_{\mathrm{BSM}} = 0$. The measured results for the combination of ATLAS and CMS are reported together with their uncertainties, as well as the individual results from each experiment. The hatched areas show the non-allowed regions for the κ_{t} parameter, which is assumed to be positive without loss of generality. The error bars indicate the 1σ (thick lines) and 2σ (thin lines) intervals. When a parameter is constrained and reaches a boundary, namely $|\kappa_{\mathrm{V}}| = 1$ or $B_{\mathrm{BSM}} = 0$, the uncertainty is not defined beyond this boundary. For those parameters with no sensitivity to the sign, only the absolute values are shown. Reproduced from [44], Fig. 15.

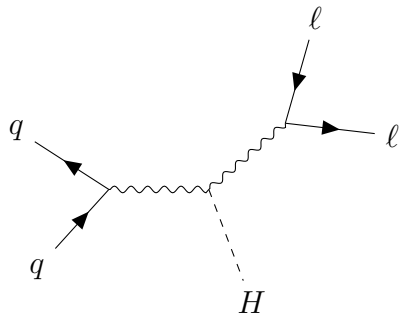

Figure 2.9: Diagram showing leading-order ZH production.

The SM predicts that a tiny 0.1% of Higgs bosons decay invisibly, through the $H \to ZZ^* \to \nu\bar{\nu}\nu\bar{\nu}$ process. The observation of a sizable branching fraction of the Higgs boson to invisible states ($\mathcal{B}(H \to \text{inv.})$) would be a strong sign of BSM physics [45, 46, 47]. Indeed, many variants of the BSM models presented here can allow a sizable branching fraction for H(inv.) decays: several "Higgs portal" models [48, 49, 50] connect the SM to a DM sector via a SM Higgs boson mediator, and if the DM candidate particle is light enough, $m_{\text{DM}} < m_{\text{H}}/2$, the Higgs boson can decay to it; in the ADD model, scalar gravitons can mix with the SM Higgs boson, allowing it to oscillate into [51] or decay to [52] particles that propagate in the extra dimensions and thus escape detection; and scalar unparticles can mix with the SM Higgs boson, which would significantly influence its properties [53]. In addition to the models presented here, countless other models that attempt to explain EWSB will have to account for the observation of a SM Higgs boson, and those which have a scalar field must allow the possibility of it mixing with the Higgs field, altering the properties of the SM Higgs boson. Given the ample opportunity for H(inv.) decays in BSM theories, and given that there are experimental methods to directly observe such decays, searches for this process are of particularly high interest.

2.6.2 Description of Higgs invisible models

In general, there is no need to introduce any particular model for H(inv.) decay, as the production mechanism is assumed to be SM. For Higgs-portal DM interpretations, it is useful to construct an explicit model of the DM sector and its coupling to the Higgs field, so that cross sections of processes involving virtual H exchange can be computed, e.g. for recasting limits on $\mathcal{B}(H \to \text{inv.})$ to limits on DM-nucleon cross sections. In [48], three models are presented, assuming scalar, fermion, and vector

($\chi = S, f, V$, respectively) DM candidates:

$$\begin{aligned}
\mathscr{L}_S &= -\frac{m_S^2}{2}S^2 - \frac{\lambda_S}{4}S^4 - \frac{\lambda_{hSS}}{4}H^\dagger H S^2, \\
\mathscr{L}_f &= -\frac{m_f^2}{2}\bar{f}f - \frac{\lambda_{hff}}{4}H^\dagger H \bar{f}f, \\
\mathscr{L}_V &= \frac{m_V^2}{2}V_\mu V^\mu + \frac{\lambda_V}{4}(V_\mu V^\mu)^2 + \frac{\lambda_{hVV}}{4}H^\dagger H V_\mu V^\mu,
\end{aligned} \tag{2.9}$$

where m_χ is the DM mass, and $\lambda_{h\chi\chi}$ the DM-Higgs coupling constant. For our purposes, the scalar and vector self-coupling terms (λ_S, λ_V) can be ignored. The fermion DM model is non-renormalizable due to the λ_{hff} term, so it is an EFT model. The vector DM model requires an additional Higgs mechanism to give the vector DM a mass [49], and therefore is disfavored. Since DM-nucleon cross section limits are typically presented as a function of DM mass, the remaining free parameter $\lambda_{h\chi\chi}$ is fixed by the assumed value of $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$, allowing the DM-nucleon cross section $\sigma_{\chi N}$ to be calculated [48]:

$$\begin{aligned}
\sigma_{SN} &= \Gamma_{\mathrm{inv.}}\frac{4m_n^4 f_N^2}{\nu^2\beta_S m_h^3(m_S+m_n)^2}, \\
\sigma_{fN} &= \Gamma_{\mathrm{inv.}}\frac{8m_f^2 m_n^4 f_N^2}{\nu^2\beta_f^3 m_h^5(m_f+m_n)^2}, \\
\sigma_{VN} &= \Gamma_{\mathrm{inv.}}\frac{16m_V^4 m_n^4 f_N^2}{\nu^2\beta_V m_h^7(m_V+m_n)^2}\left(1 - 4\frac{m_V^2}{m_h^2} + 12\frac{m_V^4}{m_h^4}\right)^{-1}, \\
\Gamma_{\mathrm{inv.}} &= \frac{\mathcal{B}(\mathrm{H}\to\mathrm{inv.})}{1-\mathcal{B}(\mathrm{H}\to\mathrm{inv.})}\Gamma_{\mathrm{SM}}, \\
\beta_\chi &= \sqrt{1-4m_\chi^2/m_h^2},
\end{aligned} \tag{2.10}$$

where $\nu = 246\,\mathrm{GeV}$ is the SM Higgs field vacuum expectation value [9], $\Gamma_{\mathrm{SM}} = 4.07\,\mathrm{MeV}$ is the SM Higgs boson width [54], and f_N is the Higgs-nucleon coupling. The latter constant represents Higgs-nucleon interaction in a low-energy regime, where g_s is no longer perturbative, and is calculated using lattice QCD to be $f_N = 0.308(18)$ [50].

Chapter 3

Previous experimental results

3.1 Introduction

This result is not the first to probe the $\mathrm{Z}(\ell\ell) + p_{\mathrm{T}}^{\mathrm{miss}}$ final state. There have been collider experiment results for this specific final state from LEP [55], the Tevatron [56], as well as previous results from LHC experiments [57, 58]. More recently, an analysis by the CMS Collaboration in the same topology using a data set with $\mathcal{L} = 2.3\,\mathrm{fb}^{-1}$ of proton collisions collected in 2015 at $\sqrt{s} = 13\,\mathrm{TeV}$ was performed [59]. Compared to the previous result, this analysis includes several significant changes: a factor of ten increase in integrated luminosity; a new technique for estimating irreducible background, which was not viable with the previous data set due to its size; improvements in the event selection, which significantly lowers the reducible background contribution; and an expanded set of BSM interpretations. Simultaneous to this result, two other CMS results targeting this final state with the same data set were released [60, 61], however both of these analyses target resonant signal models (i.e. they predict a peak in the distribution of some kinematic quantity), and therefore have significantly different analysis techniques.

In this chapter, a summary of prior experimental results is presented. As all new physics models examined here lack significant direct evidence, results are reported in terms of exclusion limits on model parameters. Exclusions are given at 95% confidence level (CL) unless explicitly specified otherwise, and represent the threshold value of a given model parameter for which the 95% CL upper limit on signal strength (the ratio of observed to expected signal cross section $\sigma/\sigma_{\text{exp}}$) crosses unity. Model parameters either above or below the threshold value are then said to be *excluded*, with the direction of exclusion dictated by the sign of the parameter's correlation with the signal cross section. In most cases, the CL_s upper limit is used rather than a confidence interval, both of which will be discussed further in Section 7.8.

3.2 Particle dark matter models

The most sensitive collider search channel for DM (in the model presented here) is the mono-Jet final state, where the leading order diagram is as in Fig. 2.6 (left), but with an ISR gluon rather than a Z boson. The second most sensitive channel is the mono-Photon final state, where the ISR Z boson is replaced with a photon. The mono-Z final state has a relatively low production rate due to the mass of the Z boson, yielding the third most sensitive result for this model where the DM mediator is assumed to only couple to quarks. However, alternative DM models involving mediator coupling to the gauge sector can make the mono-Z channel the most sensitive [62] channel for detection of DM at colliders. Future analyses in this final state will include an interpretation of the results in the context of these alternative models.

As mentioned in Section 2.4.2, the LHC run 1 ($\sqrt{s} = 8\,\text{TeV}$) CMS [29] and ATLAS [30] mono-Jet analyses probed DM through EFT operators, and hence the limits reported there require some alteration to be interpreted in this simplified model

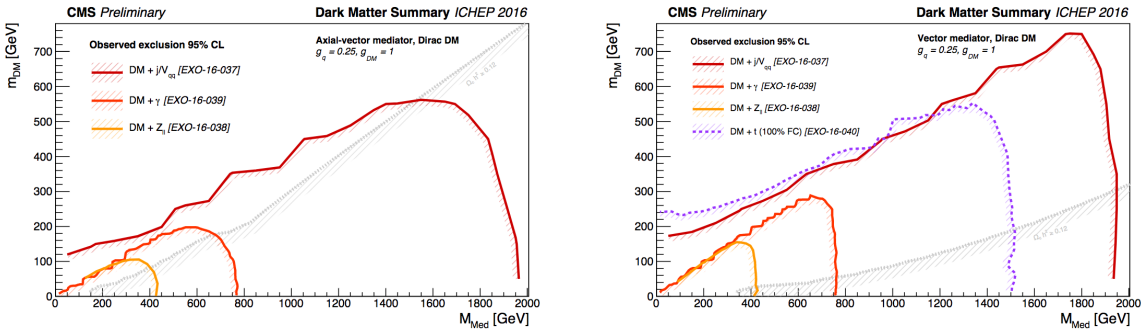

Figure 3.1: Summary plot of the various DM exclusion limits in the mediator-DM mass plane for CMS preliminary results from early 2016 data. Results are shown for mono-Jet, mono-Photon, and mono-Z channels. For vector couplings, an additional mono-Top analysis [66] is also shown, for a model where the DM-quark coupling is allowed to be flavor-changing. Reproduced from [67], Fig. 2.

context. There are also 8 TeV results in the mono-Photon [63, 64] and mono-Z [57, 58] final states from both collaborations, again using EFT models for the SM-DM interactions.

For the LHC run 2, following the ATLAS–CMS Dark Matter Forum recommendations [33], collider experiments have harmonized the reporting of DM results to use the simplified models described in Section 2.4.2. Although preliminary results are available for all channels (see Fig. 3.1), most results appearing now are based on the data collected in 2016. The most stringent limits placed so far on the simplified model parameters can be found in the $12.9\,\mathrm{fb}^{-1}$ result of CMS in the mono-Jet channel [65], which establishes lower bounds on the mediator mass of $m_{\mathrm{med}} > 1.95\,\mathrm{TeV}$ for the vector and axial vector mediated DM production models, in the on-shell ($2m_{\mathrm{DM}} \leq m_{\mathrm{med}}$) regime. For scalar and pseudoscalar mediated models, these results are not yet sensitive enough to exclude any parameter space.

As discussed in Section 2.4.2, collider searches for DM are complemented by direct detection (DD) and indirect detection (ID) experiments. DD experiments utilize a large variety of techniques, but typically involve a large sensitive volume of

some fluid or crystal, which is instrumented to detect keV-level energy deposits from DM-nucleon scattering. ID experiments search for DM annihilation via cosmic ray and neutrino telescopes. With some assumptions required, all three search strategies can place unified limits on the DM-nucleon cross section–DM mass parameter space. Scattering cross sections in DD experiments are strongly dependent on the spin structure of the DM-nucleon coupling: spin-independent cross sections increase with the atomic mass of the sensitive material; whereas spin-dependent coupling strength depends only on the nuclear angular momentum of the material. Because of this, DD searches are less sensitive to spin-dependent couplings. ID and collider searches are less affected by the spin structure, although they are less sensitive to spin-independent couplings than DD. A summary plot of several DD experimental results for spin-independent coupling is shown in Fig. 3.2. For spin-dependent coupling, ID experiments probing DM annihilation into neutrinos such as Super-Kamiokande [68] and IceCube [69, 70] are competitive with DD experiments, and place 90% CL upper limits on DM-nucleon cross sections as low as $10^{-40}\,\mathrm{cm}^2$ for DM masses in the $10 - 100\,\mathrm{GeV}$ range. The PICO-60 [71] result places the most stringent non-collider limits on the spin-dependent parameter space to date, excluding DM-nucleon cross sections down to $3.4 \times 10^{-41}\,\mathrm{cm}^2$ for a $30\,\mathrm{GeV}$ WIMP.

3.3 Large extra dimension and unparticle models

Searches for large extra dimensions have been carried out at LEP [74, 75, 76] and at the Tevatron [77, 78, 79], as well as at the LHC. The $8\,\mathrm{TeV}$ LHC [29, 30] monoJet results provide the most powerful limits on the ADD model so far, and have excluded M_D up to $5.5\,\mathrm{TeV}$ depending on the choice of $n = 2 - 6$ extra dimensions, as summarized in Fig. 3.3.

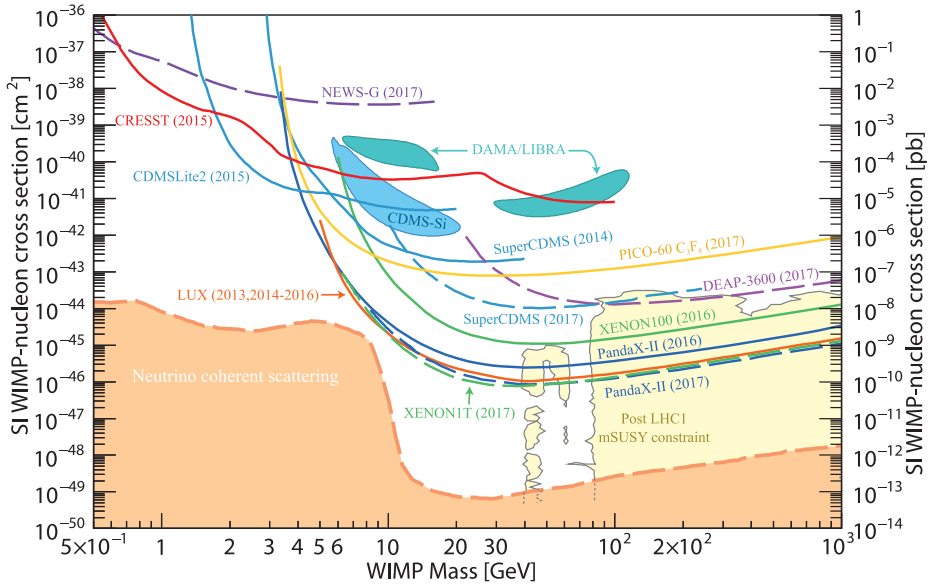

Figure 3.2: 90% CL upper limits on WIMP cross sections (normalized to a single nucleon) for spin-independent coupling versus mass. The two enclosed areas near the center represent regions of interest from possible signal events in the DAMA/LIBRA [72] and CDMS-Si [73] experiments. The beige region represents remaining allowed parameter space for selected models of supersymmetry that remain after the exclusions of LHC run 1. The orange exclusion at the bottom of the plot represents the cross section threshold at which neutrino-nucleon scattering will be detectable. For the remaining exclusion curves, the allowed parameter space is that which is below the curve. Reproduced from [9], Fig. 27.1.

For unparticles, the large choice of operator terms which can be added to the SM Lagrangian complicates the comparison of prior results, as each search channel investigates a different operator. The CMS 8 TeV mono-Jet result [29] provides an unparticle interpretation using the operator of Eqn. 2.8 with $\lambda = 1$, and places lower limits on $\Lambda_{\mathcal{U}}$ from 2 to 10 TeV for $d_{\mathcal{U}}$ between 1.8 and 1.5, respectively. The BaBar collaboration searched for the rare decay $B \to K^* \nu\bar{\nu}$, where, following theoretical motivations given in [80], a hypothetical excess could be attributed to B decays to scalar or vector unparticles. The SM branching fraction for this decay was too low to be observed [81, 82], however an upper limit of $\Gamma/\Gamma_{\text{tot}} \leq 5 \times 10^{-5}$ was established. This limit can be translated to a limit on the coupling of the pseudoscalar unparticle

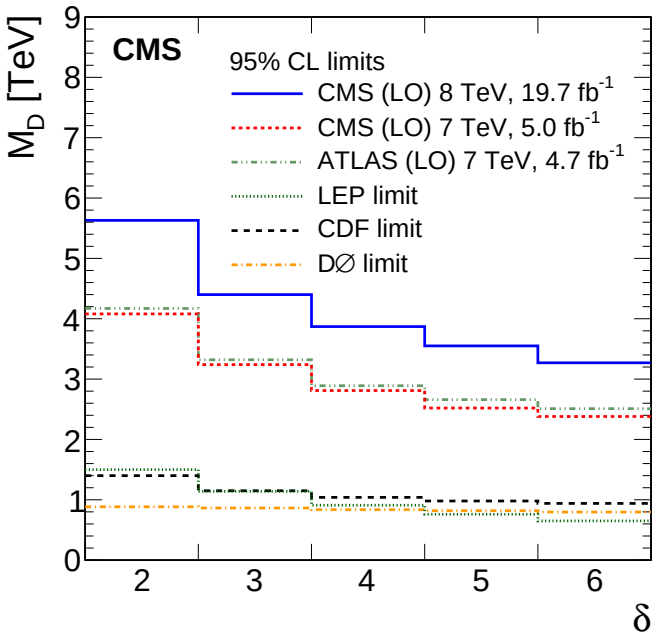

Figure 3.3: Lower limits at 95% CL on M_D as a function of the number of extra dimensions δ, obtained in the jet+$p_{\mathrm{T}}^{\mathrm{miss}}$ channel, with previous results from LEP [74, 75, 76], the Tevatron [77, 78, 79], and earlier LHC searches shown for comparison. Reproduced from [29], Fig. 7.

operator

$$\frac{\lambda_P}{\Lambda_{\mathcal{U}}^{d_{\mathcal{U}}}} \bar{s}\gamma_\mu\gamma_5 b\partial^\mu O_{\mathcal{U}},$$

which does not match the operator of Eqn. 2.8, and hence cannot directly be compared to the model investigated here. The energy scale of interactions responsible for hypothetical vector unparticle couplings to electrons has been constrained to over 4×10^5 TeV for $d_{\mathcal{U}} = 1.5$ by measurements of electron $g-2$ (anomalous magnetic moment) and positronium decays [83]. Constraints on the scale $\Lambda_{\mathcal{U}}$ of vector unparticle interactions from the invisible decay width of the Z boson and Υ meson were calculated in [84], and vary from the Planck scale to the GeV scale, for $d_{\mathcal{U}}$ from 1.5 to 2, respectively.

3.4 Invisible Higgs boson decay models

Invisible decays of the Higgs boson were searched for long before the Higgs boson was discovered. In fact, the initial 1982 paper by Shrock and Suzuki [85] in support of searches for H(inv.) was motivated by the possibility that the Higgs boson could decay predominantly to invisible final states. A reinterpretation of results from the first run of LEP [86] placed limits on the invisible branching fraction $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ as a function of the Higgs boson mass, with varied assumptions for the coupling between the Higgs and Z bosons. For $m_{\mathrm{H}} = 50\,\mathrm{GeV}$, the branching fraction was constrained below 89%, assuming SM rates of ZH associated production at this particular m_{H}. During the second run of LEP, direct searches for H(inv.) through the $\mathrm{Z}(q\bar{q})\mathrm{H}(\mathrm{inv.})$ and $\mathrm{Z}(\ell\ell)\mathrm{H}(\mathrm{inv.})$ channels were performed [87, 55, 88], probing larger values of m_{H}, and placed upper limits of $\sigma \times \mathcal{B}(\mathrm{H} \to \mathrm{inv.}) < 100\,\mathrm{fb}$ for $m_{\mathrm{H}} = 115\,\mathrm{GeV}$, for which the SM production cross section is only 50 fb.

LEP was very close to having the required beam energy to discover the Higgs boson, with a peak $\sqrt{s} = 208.9\,\mathrm{GeV}$ [89]. The threshold for on-shell ZH production is $m_{\mathrm{Z}} + m_{\mathrm{H}} = 216\,\mathrm{GeV}$. Interestingly, the 2.5 GeV Z width allows a virtual $\mathrm{Z}^*\mathrm{H}$ production at LEP-accessible energies, and a calculation in PYTHIA (see Chapter 5) gives cross sections of $1.78, 2.35, 2.98\,\mathrm{fb}^{-1}$ at $\sqrt{s} = 204.8, 206.4, 208.0\,\mathrm{GeV}$, respectively. For the integrated luminosities presented in Fig. 11 of [89], there is a 40% chance that at least one Higgs boson was produced at LEP in 2000.

Since LEP could not reach the necessary $\sqrt{s}$ to allow on-shell ZH production, these constraints were obviated after the SM Higgs boson was found at the LHC. Now that the SM Higgs boson has been identified, results can be interpreted in terms of the fraction of Higgs bosons decaying invisibly, under the assumption that the

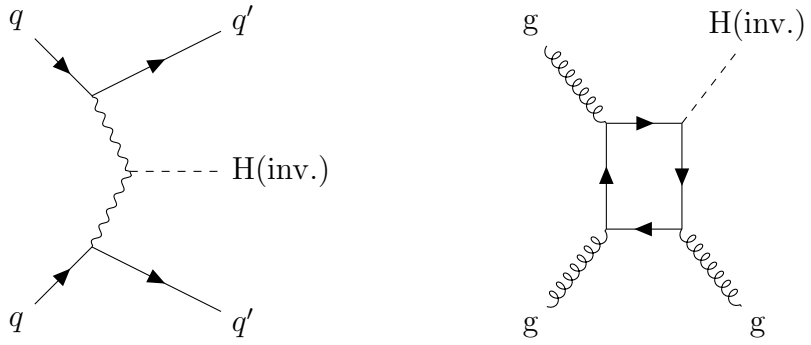

Figure 3.4: Feynman diagrams showing leading-order processes contributing to the VBF H(inv.) (left) and gH(inv.) (right) channels. The gH(inv.) process produces an identical final state as the mono-Jet process.

production cross section for the Higgs boson is that of the SM, σ_{SM},

$$\frac{\sigma \times \mathcal{B}(\mathrm{H} \to \mathrm{inv.})}{\sigma_{\mathrm{SM}}}.$$

In many contexts, this formula and $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ are used interchangeably, however it is important to remember the assumption that production follows the SM prediction. The Z($\ell\ell$)H(inv.) final state was investigated [56] at the Tevatron, where for the SM Higgs boson, the observed upper limit on the H(inv.) cross section was found to be just slightly above the SM expectation, so $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ was not constrained.

At the LHC, the most powerful channel for constraining $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ is the vector boson fusion (VBF) process. The Z($\ell\ell$)H(inv.) channel is slightly less powerful, followed by gH(inv.), which is an identical signature as the mono-Jet channel for DM. The LO diagrams for the VBF and gH(inv.) processes are shown in Fig. 3.4. The most stringent constraints available on $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ were obtained from a combination of searches in these final states at $\sqrt{s} = 8\,\mathrm{TeV}$ by ATLAS [90] and at multiple center-of-mass energies by CMS [91], which, under the assumption of SM production, exclude a branching fraction for H(inv.) decays larger than 25% for ATLAS and 24% for CMS. The relative strength of each channel in establishing this exclusion can be seen in Fig. 3.5.

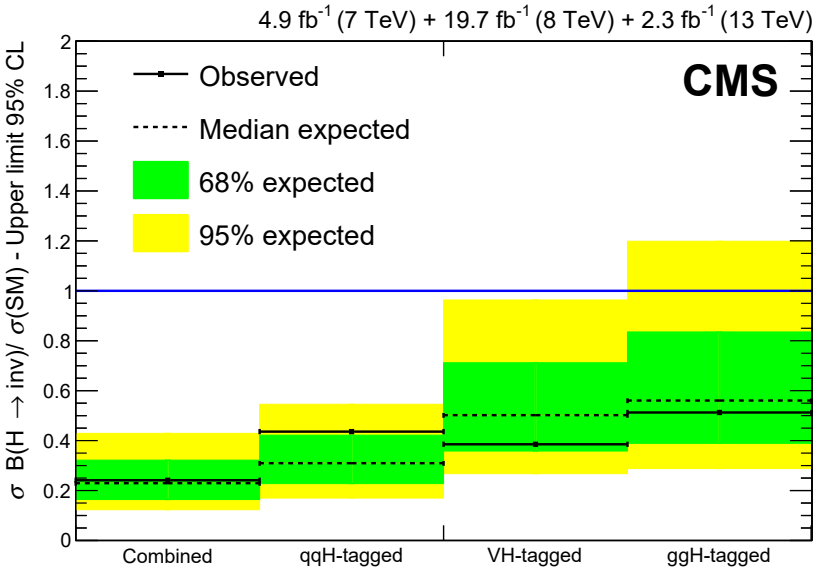

Figure 3.5: Observed and expected 95% CL limits on $\sigma \times \mathcal{B}(\mathrm{H} \to \mathrm{inv.})/\sigma_{\mathrm{SM}}$ for individual combinations of categories targeting VBF, ZH, and gH production, and the full combination assuming a Higgs boson with a mass of 125 GeV. The VBF process contributes to the qqH-tagged category. Both $\mathrm{Z}(\ell\ell)\mathrm{H}$ and $\mathrm{V}(q\bar{q})\mathrm{H}$ (where V = W or Z) contribute to the VH-tagged category, where the former is the dominant contribution to the strength of the exclusion. The ggH-tagged category contains the mono-Jet gH final state. Reproduced from [91], Fig. 6.

Chapter 4

Experiment

4.1 The Large Hadron Collider

The Large Hadron Collider (LHC) [92] at the European Organization for Nuclear Research (CERN) is the world's largest and most powerful particle accelerator and collider. It is a synchrotron situated in a 26.7 km-circumference, approximately circular tunnel that lies between 45 and 170 meters underneath the France-Switzerland border, near Geneva, Switzerland. The tunnel was constructed between 1984 and 1989 to be used for a prior accelerator/collider, the Large Electron-Positron collider (LEP), which was in operation from 1989 to 2000. The tunnel has 8 arc sections and 8 straight sections, with the arcs having a bending radius of $R = 2.804\,\text{km}$. The LHC collides protons and heavy ions in four interaction points with a beam energy of up to 6.5 TeV. The LHC does not accelerate protons from rest to high energy, but rather is fed by a sequence of smaller accelerators. A schematic of the CERN accelerator complex is shown in Fig. 4.1. The proton acceleration chain is as follows: protons from ionized Hydrogen gas are injected into a linear accelerator, Linac2, which accelerates them to a beam energy of 50 MeV; after which they are injected

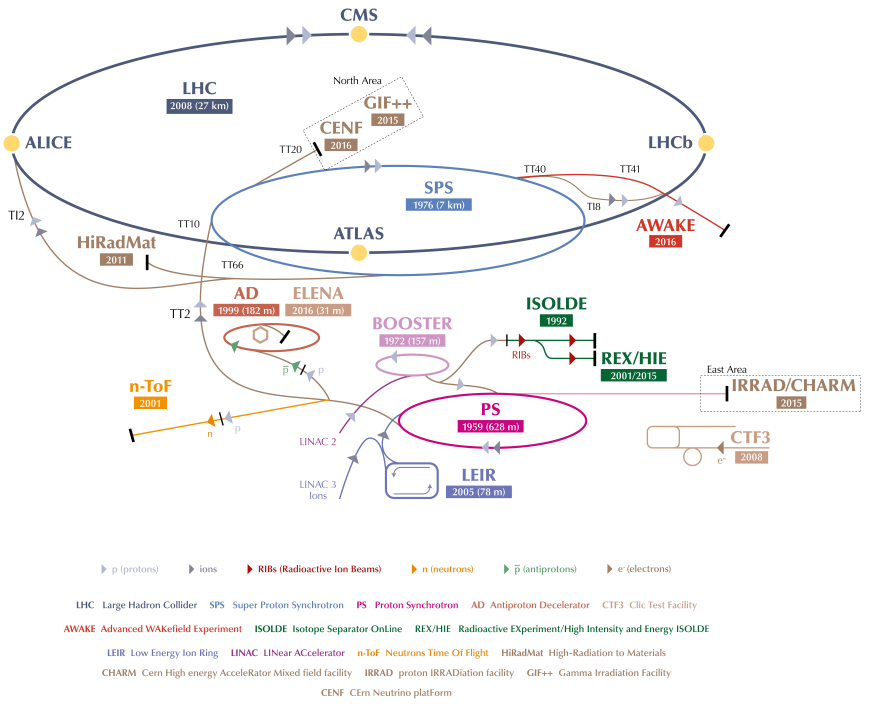

Figure 4.1: Schematic of the CERN accelerator complex, showing accelerators that feed the LHC, as well as those used for other purposes. The acceleration chain that feeds the LHC includes Linac2, the PS Booster, the Proton Synchrotron, and the Super Proton Synchrotron.

into one of the four rings of the Proton Synchrotron (PS) booster, which accelerates the protons further to 1.4 GeV; then the beams from the four rings are extracted to fill the PS ring, which accelerates the protons to 25 GeV; followed by the Super Proton Synchrotron (SPS), which accelerates protons to 450 GeV, after which the protons are sufficiently prepared for insertion into the LHC.

The ultimate beam energy of the LHC is constrained primarily by the radius of curvature R of the beamline and the strength of the bending magnetic field B,

$$E_{\text{beam}} = eBRv \approx eBRc \tag{4.1}$$

where $v \approx c$ is the speed of the beam, and e is the proton electric charge. For the

arcs of the LHC tunnel, $R = 2.804$km, leading to the design beam energy of 7 TeV if the dipole bending magnets operate at 8.33 T. It is challenging for the dipoles to reach this field strength for several reasons: at high field, the mechanical stresses on the superconducting windings of the magnet are significant, requiring high rigidity to maintain a uniform field; the self-field created by currents in the winding lowers the effective critical current of the superconductor [93]; and the magnet self-interaction stresses can cause winding movements, generating heat from friction, which can cause the wire to become resistive, leading to runaway resistive heating that can damage the magnet, known as a quench. Gradual deliberate increases of the magnetic field can be made, where controlled quenches are allowed to occur, after which the mechanical rigidity of the winding improves, and quenches no longer occur at a given field strength. This process is known as "training," and by 2015, all 1232 dipoles of the LHC became sufficiently trained to run at a field equivalent to a beam energy of 6.5 TeV [94]. An additional source of magnet quenches can originate from the heat load from protons that drift away from the beamline and impact the magnets.

There are 8 528 m-long straight sections in the LHC tunnel, which are used as insertion points in the beamline. Each of the sections has a shaft to the surface, and four of the 8 insertion points house the four main LHC experiments: A Toroidal LHC AparatuS (ATLAS) [95], at point 1; the LHC Beauty (LHCB) [96] experiment at point 8; A Large Ion Collider Experiment (ALICE) [97] at point 2; and the Compact Muon Solenoid (CMS) [98] experiment at point 5. Injection lines carrying protons from the SPS to the LHC enter in the straight sections of points 2 and 8. The four remaining insertion points contain portions of the LHC machine, including two sections of beam collimators, the beam dump, and the radio-frequency (RF) accelerator cavities.

The RF acceleration mechanism used by the LHC (and also its feeder accelerators) naturally coalesces the protons into bunches along the beamline, and the LHC bunch

spacing of 24.95 ns is formed due to the way the LHC captures bunches from the SPS in one of every 10 RF "buckets" which have a period corresponding to the LHC RF frequency of $f_{\mathrm{RF}} = 400.790\,\mathrm{MHz}$. For a successful capture into a single bucket, the beam leaving the SPS must have a narrow profile along the beam axis, which is quantified by the longitudinal emittance [99]. The total circumference of the LHC beamline corresponds to the 35640th harmonic of the RF frequency, giving 3564 possible locations for bunches and an orbit frequency $f_{\mathrm{rev}} = 11.245\,\mathrm{kHz}$. The maximum number of bunches that can be placed in the LHC is 2808, due to the necessity for gaps in the bunch trains to allow sufficient turn-on time for the injection and beam dump redirector (kicker) magnets.

4.2 LHC Operating characteristics

A key figure of merit for LHC performance is the instantaneous luminosity $\mathcal{L}$. The two proton beams are made of approximately Gaussian distributed bunches of protons that collide at the interaction points with a small crossing angle, to prevent unwanted collisions between neighboring bunches. The instantaneous luminosity can be parameterized as:

$$\mathcal{L} = \frac{N_b^2 n_b f_{\mathrm{rev}} \gamma_r}{4\pi\epsilon_n \beta^*} \left(1 + \frac{\theta_c^2 \sigma_z^2 \gamma_r}{4\epsilon_n \beta^*}\right)^{-1/2} \tag{4.2}$$

where N_b is the number of protons per bunch, n_b the number of bunches per beam, f_{rev} is the orbit frequency of the beams, γ_r the relativistic gamma factor, ϵ_n the normalized transverse emittance, β^* the beta function at the collision point, θ_c the crossing angle, and σ_z the RMS longitudinal bunch length. All of these parameters are optimized to arrive at the highest possible instantaneous luminosity, and the LHC machine operators continuously improve the performance of the machine. The original LHC design luminosity of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ was surpassed during 2016, reaching

Beam parameter	LHC design value	2016 operation value
Protons per bunch N_b	1.15×10^{11}	1.25×10^{11}
Bunches per fill n_b	2808	2076
Orbit frequency f_{rev}	11.245 kHz	
Gamma factor $\gamma_r = E_{\text{beam}}/m_p$	7463	6930
Normalized transverse emittance ϵ_n	3.75 μm	3.4 μm
β^* parameter	55 cm	40 cm
Crossing angle θ_c	140 μrad	185 μrad
RMS longitudinal bunch length σ_z	7.87 cm	
Instantaneous luminosity $\mathcal{L}$	$1.1 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$	$1.3 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$

Table 4.1: Values of the various beam parameters that enter into the instantaneous luminosity formula of Eqn. 4.2. Values are given for the original design [100] the LHC, as well as the typical operating parameters during the 2016 data-taking run [101].

a peak value of $1.5 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$. A summary of the numerical values of these beam parameters for the initial LHC design, as well as typical values attained during 2016 operation, is given in Table 4.1.

The values of Table 4.1 represent the beam conditions at the start of collisions. After the machine is filled and begins colliding protons, N_b is steadily reduced due in majority to the collisions themselves, with secondary losses in collimators and arcs. This process is exponential, with a characteristic luminosity lifetime dependent on the total inelastic cross section for *p*-*p* interactions ($\sigma_{\text{inel}} \approx 80\,\text{mb}$ [102]), which is of order ten hours. Since the collisions occur in bunches, and since the per-bunch instantaneous luminosity is sufficiently large, each bunch crossing contains multiple collisions, known as *pileup*, which is Poisson distributed with mean

$$N_{\text{pu}} = \frac{\mathcal{L}\sigma_{\text{inel}}}{n_b f_{\text{rev}}}.$$

In 2016, the peak N_{pu} was 38 collisions per bunch crossing, for the beginning-of-fill parameters of Table 4.1, and the average N_{pu} was approximately 23 collisions per bunch crossing.

4.3 The Compact Muon Solenoid experiment

The Compact Muon Solenoid (CMS) detector [98] is located approximately 100 meters underground, centered at interaction point 5 of the LHC ring. The CMS detector was designed to be a general purpose detector that would provide insight into the nature of electroweak symmetry breaking, for which the Higgs mechanism was the most likely explanation at the time. Thus, some key decisions in the design of the CMS detector were informed by its expected ability to reconstruct certain final states of Higgs boson production. However, the main requirements are not incompatible with a broad physics program, and can be summarized as follows:

- Good muon identification and momentum resolution, such that a dimuon momentum resolution of better than 1% at 100 GeV can be achieved, as well as accurate muon charge assignment up to the TeV scale;
- Good charged particle momentum resolution and reconstruction efficiency for particles with momentum as low as 1 GeV, and efficient tagging of τ leptons and b-jets;
- Good electromagnetic energy resolution such that a diphoton or dielectron mass resolution of better than 1% at 100 GeV can be achieved, while also maintaining sufficient π^0 rejection and efficient photon and lepton isolation;
- Good missing transverse momentum and dijet mass resolution.

An overview of the design of the components of the CMS detector such that it meets these requirements will be discussed in the following sections. A diagram of the detector components is shown in Fig. 4.2. The components are arranged in approximately cylindrical layers, and include, starting from the interaction point, the beam

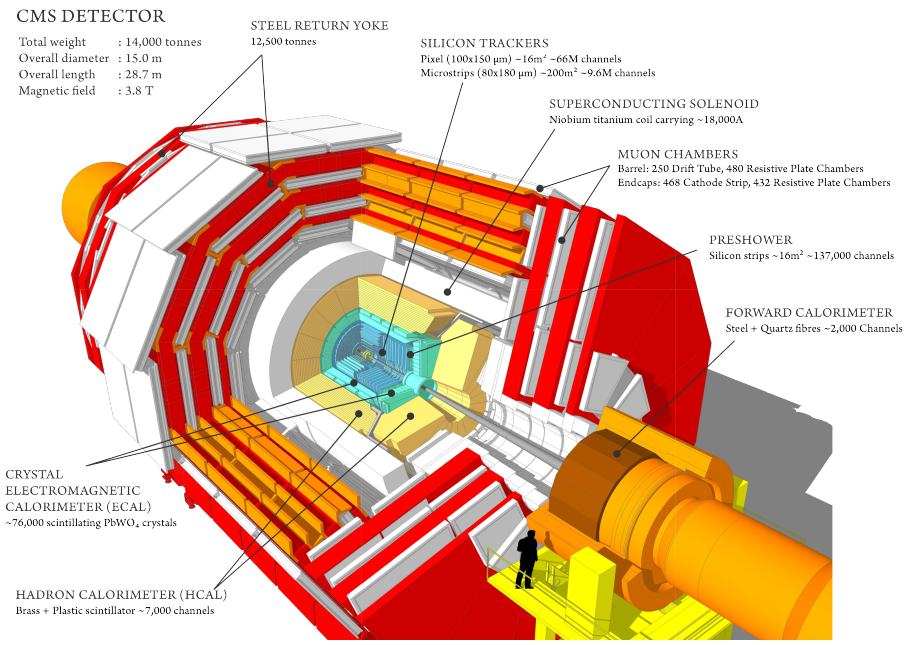

Figure 4.2: Diagram of the components of the CMS detector.

pipe, pixel detector, inner tracker, electromagnetic calorimeter, hadronic calorimeter, magnet, and muon systems, respectively.

The coordinate system adopted by CMS has the origin centered at the nominal collision point, with the y axis pointing upwards and the x axis pointing radially inward toward the center of the LHC ring. The z axis is then given by right hand rule, and points along the beamline in the direction of the Jura mountains. The azimuthal angle ϕ is measured from the x axis in the x-y plane, and the radial coordinate is denoted by r. The polar angle θ is measured from the z axis, although for the majority of applications, it is more physically meaningful to encode this degree of freedom as the pseudorapidity variable,

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

due to the fact that differences in η are invariant under Lorentz transformations along the z axis. The 3-momentum of a particle exiting the collision point is most

commonly expressed as

$$\vec{p} = (p_{\mathrm{T}}, \eta, \phi),$$

where p_{T} is the component of the momentum in the transverse (x-y) plane.

4.4 CMS Magnet

A central feature of the CMS experiment is the 6.3 m diameter, 12.5 m long superconducting solenoid. The magnet is designed to operate at a central field strength of 4 T, although the nominal operating field was de-rated to 3.8 T during the commissioning of the detector to increase safety margins [103]. The large diameter and field strength of this magnet makes it the most powerful magnet in the world in terms of stored energy, with a full-field energy of 2.5 gigajoules.

The main purpose of the solenoid is to bend the trajectories of charged particles exiting the collision point via the Lorentz force,

$$F = \frac{q}{\gamma m}\vec{p} \times \vec{B},$$

where q is the charge of the particle, m its mass, and γ the relativistic factor. This force causes charged particles to propagate outward from the collision point in an approximately helical trajectory, where deviations from the helix are due to non-uniformity of the magnetic field and multiple scattering interactions in the detector material that the particles pass through. The magnetic field inside the solenoid is modeled in simulation to better than 0.1% accuracy over the entire inner tracker volume, as validated by measurements performed prior to installation of detector components inside the solenoid [104]. Outside the solenoid, a multi-layer steel *return yoke* captures the majority of the magnetic flux, fully saturating the steel and creating an approximately 2 T magnetic field outside the solenoid. Muon detectors (Section 4.8)

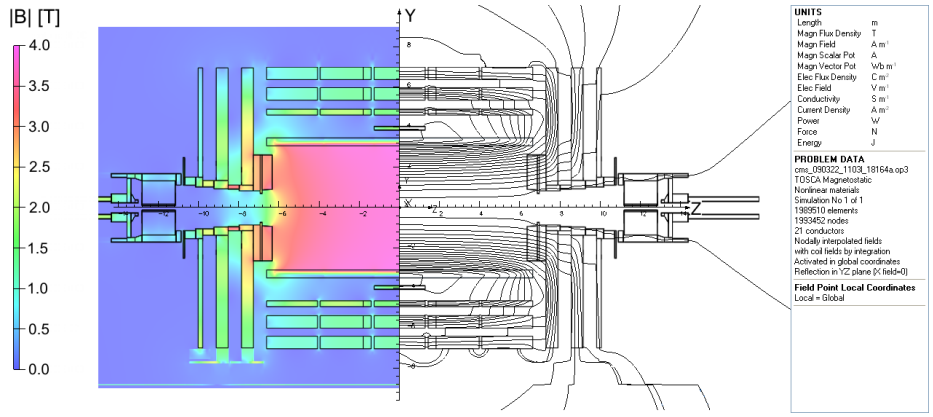

Figure 4.3: Value of B (left) and field lines (right) predicted on a longitudinal section of the CMS detector, for the underground model at a central magnetic flux density of 3.8 T. Each fieldline represents a magnetic flux increment of 6 Wb. Reproduced from [105], Fig. 5.

intersperse this structure and take advantage of the return field in the steel to measure muon momentum independently of the inner tracker. A cross-sectional view of the magnetic field is shown in Fig. 4.3.

4.5 CMS Pixel detector and inner tracker

The inner tracking system is designed to precisely and efficiently measure the helical trajectories of charged particles that emanate from the interaction point and propagate in the magnetic field created by the solenoid, which allows for the reconstruction of the charged particle momentum as well as its originating vertex. The tracker is the innermost detector element, and surrounds the beamline with a length of 5.8 m and diameter of 2.5 m, corresponding to a pseudorapidity coverage of $|\eta| \leq 2.5$. To reconstruct the particle trajectories, several layers of silicon measure small energy deposits from ionization of active material as charged particles pass through it. The minimum number of channels is dictated by the hit position resolution requirements as well as the channel occupancy limit—the probability for two charged particles to

traverse the active material of a single channel must be kept very low to prevent ambiguities in track reconstruction. With the typical maximum of 50 collisions per bunch crossing, and with about 30 charged particles produced in an average p-p collision, 1500 charged particles traverse the inner tracker every 25 ns, corresponding to a hit rate density of 1.5 MHz/mm^2 at $r = 4$ cm, decreasing at larger radii. To maintain an acceptable occupancy at this hit rate, the innermost 10 cm of the tracker uses $100 \times 150\,\mu$m^2 pixels, giving a mean occupancy of order 10^{-4}, while the inner and outer tracker (at $20 < r < 55$ cm and $55 < r < 110$ cm, respectively) use strip sensors elongated in the z direction to a typical area of 10 cm $\times\, 80\mu$m to reduce channel count (and cost) while maintaining a low hit occupancy of 2 to 3%. The pixel detector consists of 3 barrel layers and 2 endcap disks, and the strip tracker consists of 10 barrel layers and 12 endcap disks. The layers of the strip tracker are arranged such that a charged particle traverses a minimum of 9 layers in the region $|\eta| < 2.4$. The hit position resolution for this geometry is sufficient to provide track resolution better than 2% for central ($|\eta| < 1.5$) muons with momenta below 100 GeV. The complete tracker geometry is shown in Fig. 4.4.

4.6 CMS Electromagnetic calorimeter

The electromagnetic calorimeter (ECAL) lies just outside the tracker, and is designed to measure the energy of photons and electrons with percent-level accuracy. The active material consists of 61200 lead tungstate ($PbWO_4$) crystals in the barrel portion (EB, $|\eta| < 1.479$), with an additional 7324 crystals in the endcaps (EE, $1.479 < |\eta| < 3.0$.) The choice of lead tungstate was motivated by the relatively short radiation length of 0.89 cm, allowing a radially compact detector in the interior of the solenoid. In addition, the 2.2 cm Molière radius allows very precise shower

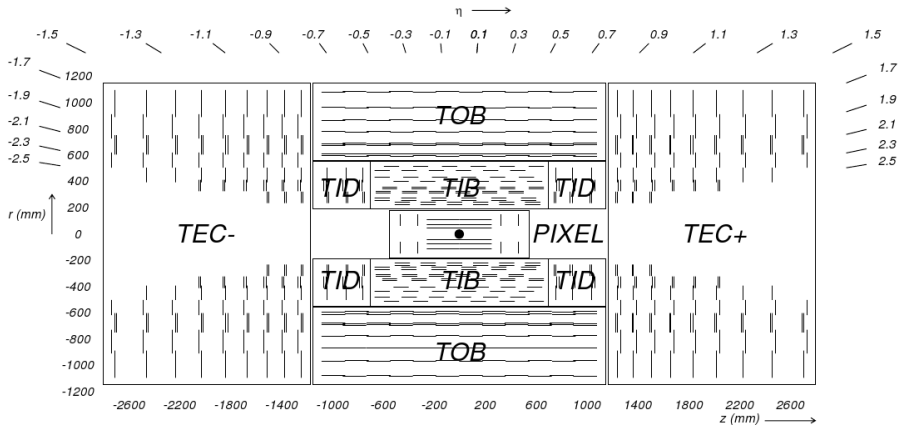

Figure 4.4: Diagram of the CMS inner tracker, comprised of the pixels, tracker inner barrel (TIB), tracker outer barrel (TOB), tracker inner disk (TID), and tracker endcap (TEC) components. Each line represents a detector module. Double lines represent stereo modules, where two two modules are mounted back-to-back to provide two closely-spaced hit layers.

position measurement, which is crucial for good electromagnetic object momentum measurement. For EB, the crystals have a truncated pyramid geometry, with a front face one Molière radius wide, and a depth corresponding to 25.8 radiation lengths. The EE crystals have a tapered rectilinear geometry, and have a slightly larger front face with sides of length 2.86 cm and a reduced depth of 24.7 radiation lengths. The barrel crystals are arranged in a semi-projective geometry, pointing slightly off-center from the interaction point to minimize gaps in material, while the endcap crystals are arranged in a rectangular grid, with the crystals pointing at a focus 1.3 m beyond the interaction point. The crystal dimensions in the η-ϕ plane are 0.0174×0.0174 in the central barrel, while in the endcap the dimensions become increasingly irregular as $|\eta|$ increases. Scintillation light is collected at the rear face of the crystals and converted to electrical pulses in EB via avalanche photodiodes (APDs), and in EE via vacuum phototriodes (VPTs.) The barrel crystals are grouped into 5×5 *towers*, and the endcap crystals are grouped to approximate projective towers, which are used by the off-detector electronics to provide trigger information (Section 4.9.1) as

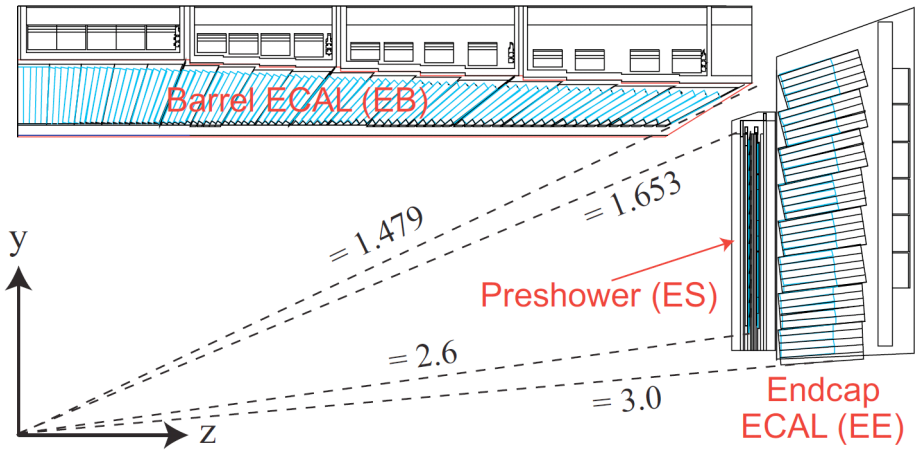

Figure 4.5: Diagram of the CMS electromangetic calorimeter, including the barrel, endcap, and preshower components. Dashed lines show reference pseudorapidity values with respect to the center of the detector.

well as to select regions of interest for which the full crystal readout is performed. Full readout of all 68524 channels for every event is not possible due to bandwidth and event data size constraints. A silicon-lead sampling calorimeter, known as the ECAL preshower, lies in front of EE, and provides improved hit resolution in the $1.653 < |\eta| < 2.6$ region to better discriminate photons against π^0s decaying into photon pairs. In addition, the preshower compensates for the reduced depth of EE with 3 radiation lengths of lead. A schematic of the ECAL geometry is shown in Fig. 4.5.

The characteristic resolution of ECAL was established with fixed-energy electron test beams, with a partial detector assembly [106]. The resolution $\sigma(E)/E$ is parameterized in three components: a stochastic term, representing shower fluctuations and photostatistics, of $2.8\%/\sqrt{E}$; a constant term, representing crystal non-uniformity and light yield, as well as intercalibration uncertainty, of 0.3%; and a noise term, representing electronic noise and pileup contributions (both in-time and out-of-time), of 41.5 MeV per crystal, with negligible correlations. The *in-situ* resolution is slightly

larger due to time variation of the above factors and the uncertainty in intercalibration measurements to correct for these effects. In particular, a cumulative transparency loss due to radiation damage in the crystals is monitored and corrected via a laser and light-emitting diode light injection system.

4.7 CMS Hadronic calorimeters

To properly measure jets and transverse momentum imbalance of a collision, it is essential to capture hadronic energy with good resolution and hermeticity. The hadronic calorimeter (HCAL), situated between the ECAL and the magnet, covers the same geometric region as ECAL, extending to $|\eta| < 3$ and divided into barrel (HB) and endcap (HE) portions at $|\eta| = 1.3$. It is a brass-scintillator sampling calorimeter, where the brass functions as a showering material, and was selected based on the requirements that it be non-magnetic and structurally rigid, with a short hadronic interaction length, and inexpensive. The depth of HCAL is equivalent to 5.8 to 10.6 hadronic interaction lengths, varying as $1/\sin\theta$, with the most central portion having the lowest material budget. The ECAL provides an additional 1.1 interaction lengths of material. The plastic scintillator is arranged into tiles which have fiber optic cables embedded into them. The fibers carry scintillation light to readout boxes installed in the HB-HE boundary area, where an optical decoding unit arranges the fibers into readout *towers* corresponding to rectangular η-ϕ regions of the detector. The fibers for each tower are fed to a hybrid photodiode [107] (HPD), which converts the scintillation light to electrical signals. There are 16 scintillator layers in HB and 18 in HE, which are combined by the optical decoding units into a smaller number (1 to 3) of readout channels (*depths*) for each tower. To compensate for the low number of interaction lengths at low pseudorapidity, an additional outer hadronic

calorimeter (HO) is installed outside the solenoid, which uses the same scintillator tile plus HPD arrangement as HCAL, sampling the showers developed in the steel of the return yoke, and increases the amount of instrumented material in the barrel to at least 11.8 hadronic interaction lengths. The forward hadronic calorimeter (HF) lies outside the muon system, 11.15 m from the interaction point on both sides, and covers $3 < |\eta| < 5$, providing a significantly enhanced hermeticity. The active material of HF consists of quartz fibers embedded in steel showering material, chosen for radiation hardness as HF will experience 10 MGy of accumulated radiation exposure after 500 fb^{-1} of accumulated luminosity. In HF, Cherenkov radiation from particles traversing the quartz fibers is fed to photomultiplier vacuum tubes inside a radiation-shielded enclosure. Two lengths of quartz fibers are embedded in the steel: long fibers collect Cherenkov light in the whole of HF and short fibers collect light in the rear of HF. This allows for differentiation between electromagnetic showers, which deposit energy primarily at the front of the calorimeter, and hadronic showers, which deposit energy throughout HF. A schematic of the hadronic calorimeter systems is shown in Fig. 4.6.

The characteristic resolution of the barrel hadronic calorimeter was established in fixed-energy pion test beams, with a partial assembly, including a section of ECAL in front of the HCAL assembly, as well as a section of HO placed behind an aluminum slab to simulate the dead material of the solenoid [108]. The resolution is parameterized in the same was as for ECAL. The stochastic and constant terms were measured to be $115\%/\sqrt{E}$ and 5.5%, respectively. The noise term is about 180 MeV per channel.

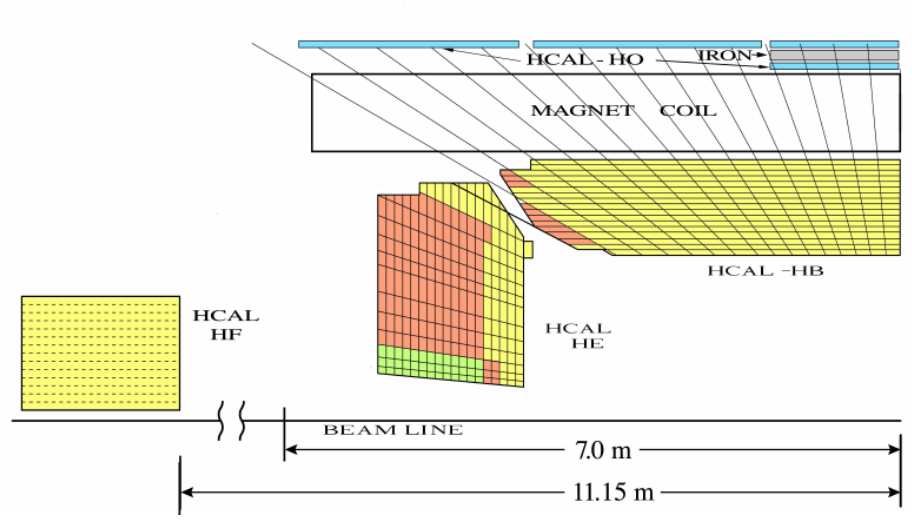

Figure 4.6: Diagram of the CMS hadronic calorimeter, including the barrel (HB), endcap (HE), outer (HO), and forward (HF) components. Shading of the layers of HB and HE show the longitudinal (depth) segmentation of the readout channels.

4.8 CMS muon detectors

Although muons leave tracks in the inner tracker, the tracker is not able to distinguish them from more commonly produced charged hadrons. To efficiently discriminate muons, the muon systems reside in the outermost layers of the CMS detector, interspersed with the steel of the return yoke, where the large number of hadronic interaction lengths (10-25, depending on the layer of the return yoke) removes the vast majority of charged hadrons[1]. Due to the desire to have as large a geometric acceptance as possible, the muon systems must cover a very large surface area, which makes cost considerations an important constraint on channel count. Simultaneously, the position resolution must also be sufficient to reach the momentum resolution goal of better than 1% for a 100 GeV muon, as well as charge assignment up to the TeV scale. Since no tracker information is fed to the trigger system, the muon systems must also provide trigger information such that events where the predominant inter-

[1] The hadrons that survive long enough to reach the muon systems are referred to as *punch-through*.

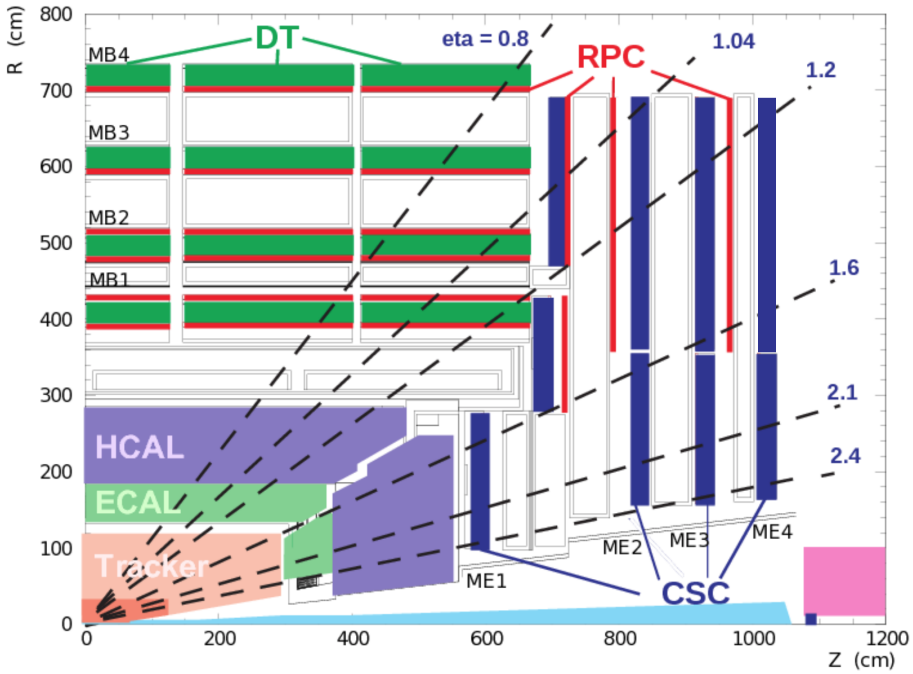

Figure 4.7: Diagram of the CMS muon systems, including the drift tubes (DTs), resistive plate chambers (RPCs), and the cathode strip chambers (CSCs). Dashed lines show reference pseudorapidity values with respect to the center of the detector.

action energy is only carried by muons can be recorded. In the endcaps, the muon systems must also accept high muon occupancies from both collision and non-collision sources. These considerations, among others, lead to the choice of three complementary muon systems: drift tubes (DTs), cathode strip chambers (CSCs), and resistive plate chambers (RPCs). A schematic overview of the muon systems is shown in Fig. 4.7.

4.8.1 Drift tubes

Drift tubes cover the barrel region of the CMS detector, providing muon detection up to $|\eta| = 1.2$. The drift tubes are rectangular tubes with a $13 \times 42\,\text{mm}^2$ cross section and a 2.4 m length. They are organized in superlayers 4 tubes deep, with each layer staggered by half a tube. Three superlayers are arranged, with alternating r-ϕ wire

alignment, into *stations*. Four layers of stations are placed, alternating with the three layers of yoke steel, in each of the segments of the dodecagonal return yoke. The stations are split into 5 segments along the z axis due to the segmentation of the yoke, leading to a channel count of about 172 000. Each tube contains an anode wire, two cathode strips, and two field-forming strips, held at specific electric potentials. Free electrons, created by ionizing radiation traversing the Ar-CO_2 gas mixture, drift in the electric field towards the anode wire, where they trigger an avalanche in the high-field region near the wire. The subsequent electrical pulse is then fed through a shaping preamplifier and digitized with a time-to-digital converter (TDC) to be read out. The TDC resolution of 265 ps gives a single-wire position resolution of 250 μm. In the barrel region, the muon rate is low and the magnetic field is mostly contained in the steel of the return yoke (see Fig. 4.3), such that relatively long drift times (up to 380 ns) and large drift distances are permissible, leading to this configuration of tube size and channel count.

4.8.2 Cathode strip chambers

Cathode strip chambers cover the endcap regions of the CMS detector, providing muon detection in the region $1.2 < |\eta| < 2.4$. The chambers are trapezoidal, with dimensions such that they cover 10° in ϕ for the inner ring of chambers, and 20° for the outer chambers. They are arranged into stations in an overlapping manner to remove dead zones, and there are 4 stations interleaved with the 3 disks of the endcap return yoke, for each side of the detector. The chambers are multiwire proportional chambers, with 6 planes of anode wires interleaved among 7 cathode strip planes. The strips run radially outward, with constant $\Delta\phi$ width corresponding to about 8 mm at the narrow end of the chamber and 16 mm at the wide end. The anode wires run along the ϕ coordinate, spaced 2.5 mm apart, and in the innermost chambers,

the wires have a slight tilt to compensate for the shift in arrival position on the wires of ionization electrons due to the magnetic field. A gas gap of 9.5 mm (7 mm in the innermost chambers) between the anode wires and cathode strips contains an Ar-CO_2-CF_4 mixture, where the ion-electron pairs are created when muons pass through. Anode wires are read out in groups of 16, and each cathode wire is read out. Anode and cathode channels are both zero-suppressed based on the presence of trigger primitives (anode/cathode local charged track) in the chamber. The larger muon rate motivates the finer granularity of the components compared to that of the DTs, and the CSC system contains about 400 000 readout channels. The small drift gap allows effective operation in the large and complex magnetic fields present in the endcaps, and the spacing of the cathode strips and anode wires is sufficiently small to achieve a 150 μm (75 μm for the inner chambers) spatial resolution.

4.8.3 Resistive plate chambers

Resistive plate chambers are present in both the barrel and endcap regions of the CMS detector, covering the region $|\eta| < 1.6$. The RPCs complement the DTs and CSCs in their respective regions by trading a reduced position resolution for the ability to localize ionization events in time to much better resolution than the 25 ns bunch spacing, with little to no dead time. This feature makes the RPCs valuable for triggering purposes. Six layers of RPCs are embedded alongside DTs in the barrel yoke, and 3 layers are installed alongside CSCs in the endcap disks. The RPC chamber is a rectangular flat sheet with a double gas gap of 2 mm separating layers of high-resistivity plastic panels. The panels are held at a high electrical potential, where an ionization in the gas gap promotes an electric discharge between them. The plates are capacitively coupled to readout strips, which capture the discharge pulse. In the barrel, the strips run along the z direction, with a size corresponding to $5/16^\circ$ in ϕ,

projectively, from the interaction point. The endcap RPC chamber is trapezoidal, and has strips running radially with the same ϕ segmentation as in the barrel. The gas is a mixture of R134a refrigerant, isobutane, and sulfur hexafluoride. The RPC channel count is about 165 000.

4.9 CMS Trigger and data acquisition

The nominal 25 ns bunch spacing leads to a bunch collision rate of 40 MHz. The total event data from all detectors, even after significant zero-suppression and compression, amounts to about a megabyte. To be able to read out, record on tape, reconstruct, and analyze this data, it must be filtered in stages to a manageable event rate. This is accomplished in stages: the level 1 (L1) trigger system selects events, based on reduced information from the muon systems and calorimeters, for readout such that the rate is reduced below 100 kHz; after which the data acquisition (DAQ) system assembles the complete event data from all detectors; which is then forwarded to the high-level trigger (HLT), which analyses the events further, and selects a subset such that the rate of events recorded to tape is kept near 1 kHz. An overview of how these systems function is presented in this section.

4.9.1 Level 1 trigger

The L1 trigger system [109, 110, 111] is designed to select events of interest based on information collected from ECAL, HCAL, and the muon systems. It assembles the information and reduces it to a set of primordial physics objects (trigger objects) corresponding to the physics objects that are reconstructed from the full detector information (Chapter 6) after event collection, where minimal data volume or processing time constraints, and no latency constraints, are present. The trigger objects

are assembled in two pipelines, as shown in Fig. 4.8: the calorimeter trigger and the muon trigger. The calorimeter trigger pipeline starts with the generation of summed transverse energies from the towers of ECAL and HCAL, along with additional quality information. Towers are the smallest unit of detector area for the calorimeters at trigger level, and the ECAL and HCAL detector channels are mapped to towers (e.g. 25 EB crystals per HB tower) such that each tower covers the same projective η-ϕ region of both calorimeters. The layer 1 calorimeter trigger combines the ECAL and HCAL tower information and sends it to one of the nine layer 2 calorimeter trigger processors, which are time-multiplexed so that each processor can process a full view of the detector. In the layer 2 processor, L1 EGamma (electron or photon) objects are formed from small, dynamically sized (up to 8 towers) clusters of towers with predominantly electromagnetic energy deposits, L1 Jets are formed from 9×9 tower energy sums, L1 Taus are formed from the same dynamic clusters as L1 EGamma objects with an energy deposit pattern consistent with hadronic decays of τ leptons, and vector and scalar energy sums over the entirety of the calorimeters are computed. These objects are then sent to the demultiplexer processor, which sorts the objects in p_{T} and forwards them to the global trigger. The muon trigger pipeline starts with the generation of track segments in DT and CSC, and strip hits in RPC. The segments and hits are then distributed via intermediate processors according to their $|\eta|$ position into the barrel, overlap, and endcap track finders. The track finder processors reconstruct muon tracks, creating a L1 Muon with charge, p_{T}, and quality information defined. The global muon trigger combines the L1 muons from the three track finders, disambiguates them, sorts them in p_{T} and quality, and delivers them to the global trigger. The global trigger tests for multiple criteria based on the presence of objects or combinations of objects with minimum energy or quality thresholds, each of which is referred to as a *trigger path*, and sends the decision to accept or

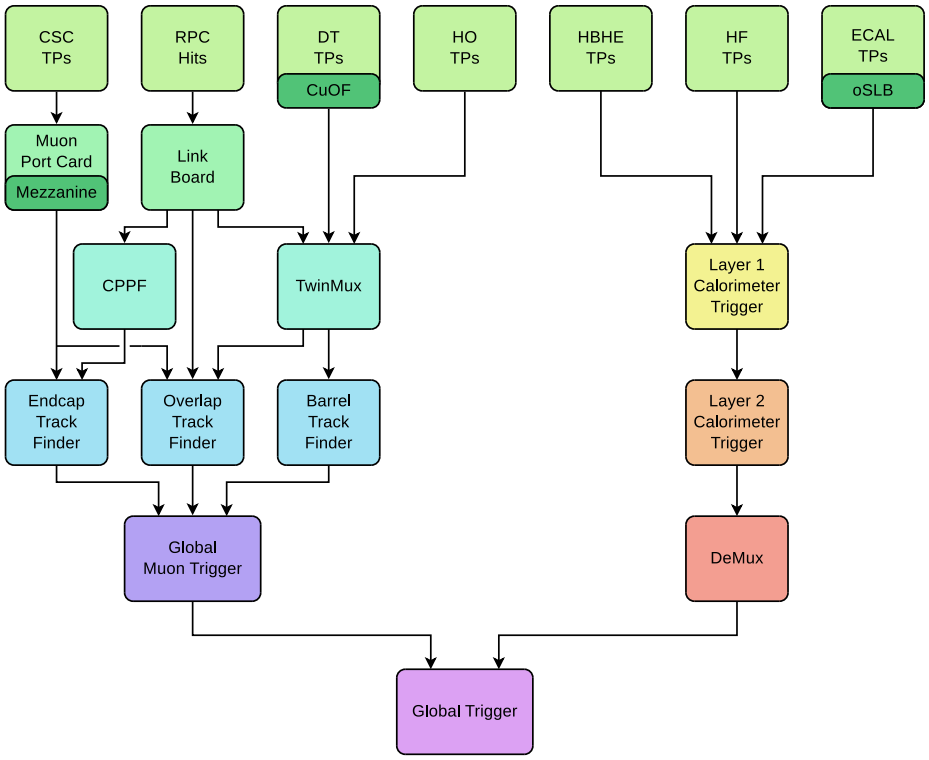

Figure 4.8: Schematic of the CMS Level-1 trigger system.

reject the event—the L1 accept—to the data acquisition system. The *trigger menu* is the collection of trigger paths, and is optimized to collect as much interesting physics data as the rate limit allows. The total trigger rate is highly correlated with the instantaneous luminosity. Some paths are *prescaled*, such that for some prescale n, only 1 in n events is recorded. These prescaled events are typically used for data cross-validation and trigger efficiency calculations.

4.9.2 Data acquisition

The data acquisition system [112] begins with the detector front-end buffers, which are often located on the detector itself. These buffers store the digitized readout of the detector channels in a pipeline of fixed length, waiting on a decision from the L1 trigger system before sending the data to the next stage of the DAQ system.

The L1 trigger must make a decision within a latency of $4\,\mu$s, corresponding to the event capacity of the smallest of these front-end buffers[2]. If the L1 trigger accepts an event, the L1 accept signal is routed to these front-ends, which then sends the data to readout devices, which format the data into packets and sends it through a high-bandwidth network switch system to the event builder devices. From the readout devices onward, the DAQ system uses commercial server and networking hardware. In this high-bandwidth (order terabits per second) builder network, the data is transformed from parallel to serial in detector components, and from serial to parallel in events.

The bandwidth limitations of the front-ends and the builder network limit the maximum rate of L1 accepts to less than 100 kHz, which is enforced by a fast feedback loop known as the trigger throttling system. The throttling system allows any component of the DAQ system to signal that its data buffers are nearing capacity, which then vetoes additional L1 accepts such that no accepted event is lost in the DAQ system. The throttling feedback loop cannot operate with sufficient latency to prevent buffer overflows from sub-microsecond bursts of triggers, due at least to the propagation delay between the components, and therefore fixed *trigger rules* prevent such bursts by vetoing more than k L1 accepts in N consecutive bunch crossings, where $(k, N) \in \{(1, 3), (2, 25), (3, 100), (4, 240)\}$. L1 accepts that are vetoed by the throttling system, trigger rules, or other reasons such as detector calibration sequences are lost, and the fraction of lost L1 accepts is known as *deadtime*, which is carefully tracked so as to properly account for it in the integrated luminosity. Typical deadtime values are around 1% for the dataset used in this analysis.

After the data has been parallelized by event, the builder units distribute events

[2] Namely, the 192 bunch crossing deep analog pipeline of the silicon strip tracker APV25 front-end ASIC. With 24 bunch crossings of latency reserved for readout operations, this leaves $168 \cdot 25\,\text{ns} = 4.2\,\mu$s of latency available for the trigger decision.

to filter units, which perform the functions of the high-level trigger. If the filter unit accepts an event, it is forwarded to a high-capacity network filesystem, where it then waits for transfer to tape at various computing sites around the world.

4.9.3 High-level trigger

The high-level trigger is a software device running on the filter units of the DAQ system. Each filter unit has access to the entire detector readout of the event in its raw form. To decide if the event should be kept, the filter unit can in principle reconstruct the physics objects of the event to the level of a final analysis. However full reconstruction the event (the subject of Chapter 6) is a computationally expensive operation, and since the number of filter units is finite (the HLT computing cluster is significantly smaller than the sum of computing resources available for post-data collection processing of the events) and the incoming event rate can be as large as 100 kHz, some compromise has to be made between the speed and quality of the HLT reconstruction. Examples include: regional reconstruction of some objects based on the positions of the L1 objects, which are precomputed by the L1 trigger and read out; seeding of tracks only by sufficiently high momentum muons or calorimeter deposits; and reduced iterations of track-finding. As in L1, the filtering is organized into trigger paths. To further reduce the decision time, reconstruction and filtering operate in turn in a given trigger path, with faster reconstruction steps placed early in the sequence, so that if an event fails one of the filter criteria with partial reconstruction, then later reconstruction steps in the trigger path need not be carried out (unless a separate trigger path requires it.) For the size of the HLT computing cluster operating in 2016, the average processing time per event was kept below 160 ms, such that the event buffers of the filter units were kept at a safe occupancy. As in L1, the menu is optimized for maximum allowable event rate, although the HLT rate is

primarily constrained by offline computing and storage cost limitations rather than by bandwidth. Some HLT trigger paths are prescaled, as in L1, for monitoring and validation purposes.

4.10 CMS Luminometers

An accurate estimate of the number of SM background events is crucial for this analysis, and a necessary ingredient for this prediction is an accurate measurement of the integrated luminosity corresponding to the collected data used in this analysis. Although in principle the instantaneous luminosity at the collision point is given by Eqn. 4.2, in practice the Gaussian distributed bunch assumption as well as the uncertainty of some parameters in this formula lead to an unacceptably high uncertainty. To accurately measure the instantaneous and integrated luminosity, CMS uses rate measurements from a variety of observables collected from five detectors: the pixel detector, DT, HF, and two detectors specialized for luminosity measurement purposes, namely the fast beam conditions monitor (BCM1f) and the pixel luminosity telescope (PLT.) [113] The observables are chosen to have a linear correlation with luminosity and stability over time, and for each, a relationship between the rate of events R and the instantaneous luminosity can be fixed by a visible cross section parameter σ_{vis},

$$\mathcal{L} = R\sigma_{\text{vis}},$$

providing a stable relative measurement of the luminosity. To establish the absolute scale, σ_{vis} must be measured in a van der Meer scan [114], where the LHC beams are scanned across each other to experimentally establish the beam profile, which when combined with accurately measured beam intensities, gives an absolute scale for the instantaneous luminosity. Values of calibrated instantaneous luminosity are made

available for each *lumisection*, which represents a fixed period of event-taking equal to 2^{18} orbits of the beams, or 23.3 s. The integrated luminosity for the dataset used in this analysis is measured with a 2.5% systematic uncertainty [113].

4.11 CMS Performance and data quality

The LHC operated in proton physics production mode from April 22 to October 27, 2016. During this time, 40.8 fb^{-1} of integrated luminosity was delivered to the CMS interaction point. Of this, 37.8 fb^{-1} was recorded by CMS, with losses due to various issues such as: temporary failures of detector components or services to them (e.g. power supplies) that prevented data collection, trigger deadtime, or human error. For all recorded data, detector experts for each detector system evaluate the quality of the data, looking for significant dead zones or spurious noise that may significantly impact the ability to analyze the data. If no issues are found, the data is certified. Data is certified or rejected in units no smaller than a a lumisection. The total integrated luminosity of the lumisections collected in 2016 that are certified good for analysis is 35.9 fb^{-1}. The cumulative integrated luminosity recorded during the 2016 run is shown in Fig. 4.9.

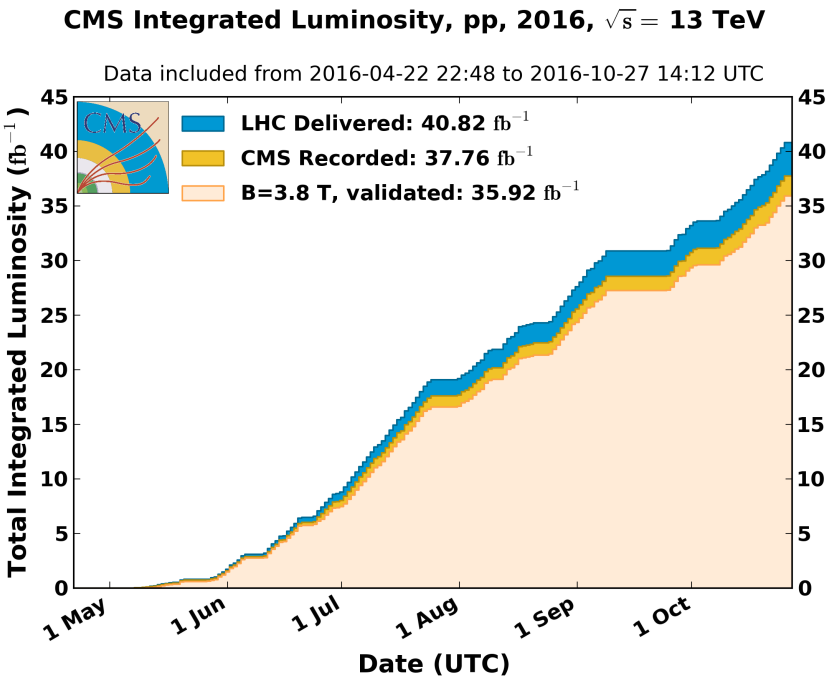

Figure 4.9: Cumulative integrated luminosity produced, recorded, and certified by the CMS experiment in 2016.

Chapter 5

Event simulation

5.1 Introduction

Accurate simulation of event data—as would be collected by the CMS experiment—is paramount to the success of this analysis. The information we want to collect from the physics processes of interest is diluted by backgrounds (discussed in Chapter 7), smeared by detector resolutions (discussed in Chapter 4), and distorted by reconstruction inefficiencies (discussed in Chapter 6). To provide accurate estimates of the background contributions, the signal selection efficiency, and the resolution of kinematic quantities, a complete simulation of the experiment is required, starting from the initial state of two protons colliding, and ending with event data in a format indistinguishable from that of the digitized detector signals that are read out from the CMS detector. Simulation is divided into several stages, each of which is explained in the following sections.

5.2 Proton-proton collision

A proton-proton collision involves the possibility of multiple interactions between the parton constituents of the proton, including the *valence* quarks that constitute the low energy scale description of the proton ($p = uud$), the gluon mediators that bind these quarks together, as well as the *sea* quarks that are spontaneously created from the gluon field, all of which have significant contributions at the energy scale of proton collisions at the LHC.

The modern picture of the proton composed of fundamental partons has emerged from the results of deep inelastic scattering (DIS) experiments over the course of several decades, beginning with the first experimental evidence [115] of Bjorken scaling [116], which established the composite nature of the proton. In these experiments, a fundamental particle (such as an electron or neutrino) is made to collide with a proton, and the scattering cross section of the incident particle is measured as a function of collision energy and scattering angle. From the results of these experiments, a probabilistic model of the proton constituents, known as a *parton distribution function* (PDF), is developed. The PDF describes the probability $f_i(x; Q^2)$ that a parton species i carries a fraction x of the total momentum of the proton, as probed by the incident particle via exchange of some gauge boson with momentum transfer Q^μ, such that the total probability, [117]

$$\sum_i \int_0^1 dx\, x f_i(x; Q^2) = 1. \tag{5.1}$$

The role of Q^2 in the f_i is to parameterize how the relative fraction of proton momentum carried by a given parton species evolves with the energy scale of the interaction, primarily as a consequence of the changing coupling g_s. Thus, the f_i must evolve with Q^2 as specified by the renormalization group equations of QCD for a given order in perturbation theory, such that the relation of Eqn. 5.1 remains con-

served. The differential equations describing this evolution are known as the DGLAP equations [118, 119, 120], where the *order* of the PDF (LO, NLO, etc.) refers to the order in QCD perturbation theory under which the DGLAP equations are derived. The f_i are simple parametric functions of x, specified at a fixed reference value Q_0^2, and evolved to arbitrary Q^2 with the help of the DGLAP equations, such that data collected by DIS experiments over several orders of magnitude in x and Q^2 can be fit to a single function.

The impressive agreement between the DGLAP-evolved $f_i(x;Q^2)$ and the DIS data taken over a wide range of x and Q^2 (see, for example, [121]) provides some of the strongest support for QCD theory, however it does not by itself describe the behavior of a proton-proton collision. The remaining ingredient is the *factorization theorem*, first introduced to predict the cross section of the Drell–Yan process [122] in which a quark from one proton and an antiquark from another proton annihilate into a virtual photon, which subsequently produces a pair of charged leptons. The factorization theorem postulates that the proton-proton collision contains two independent components: a *hard scatter* process, where two partons interact at some high energy scale; and the *underlying event*, describing the low energy scale interactions of the remnants of the protons after considering the hard scatter. Further, the cross section of a process with a given final state from the hard scatter can be found by integrating the partonic cross section of the hard scatter, weighted by the PDFs of each parton in its respective proton evaluated at some given *factorization scale* μ_F (c.f. Q of DIS),

$$\sigma_{pp\to X} = \sum_{i,j} \int dx_i dx_j \, \sigma_{ij\to X}(x_i, x_j) \cdot f_i(x_i;\mu_F^2) \cdot f_j(x_j;\mu_F^2), \tag{5.2}$$

where the X in $pp \to X$ refers to the final state created by the hard scatter, without regard to the final state of the underlying event. Similar to the renormalization scale

discussed in Section 2.1, the factorization scale is unobservable, and if one were able compute $\sigma_{pp\to X}$ to all orders in perturbation theory, it would have no dependence on μ_F. In practice, the cross section will have some μ_F dependence (albeit minimized by the use of higher order PDFs and partonic cross sections) and the choice of μ_F is usually motivated by some characteristic energy scale of the hard process, e.g. the transverse mass[1] of the final state of the hard process. In any process in which the partonic cross section calculation requires the assumption of a renormalization scale μ_R, it is usually set to the same value as μ_F, as the value of both are motivated by the scale of the process [123].

To simulate events, one must sample the phase space of the final state given by Eqn. 5.2. This is accomplished via the Monte Carlo method [124], in two stages: the matrix element (ME) generator samples the PDFs at the hard interaction scale μ_F to determine the parton species i, j, and samples the differential partonic cross section $\sigma_{ij\to X}$ to determine the kinematics of the final state of the hard process, as viewed at the hard interaction scale μ_R; then the parton shower (PS) generator determines the evolution of initial state quarks as well as any final state quarks of the hard process from the hard interaction scale to the hadronization scale, and determines the final state of the underlying event. The ME and PS generators are discussed further in the following sections. In this analysis, the NNPDF3.0 [125] PDF is used, with the order of the PDF corresponding to the order of the ME calculation used for the hard process under consideration. The PDF is evaluated with the assistance of the LHAPDF6 library [117], which interpolates or extrapolates from the values given by NNPDF, and assists with evaluation of PDF uncertainties.

[1] Transverse mass is defined $m_T^2 = E^2 - p_z^2$.

5.3 Matrix element

In this analysis, several ME generators are employed, as the order in QCD as well as the accuracy of calculation varies across generators and processes. Some generators, such as MadGraph5_aMC@NLO [126], are able to construct the matrix element for a given process at LO and NLO QCD in a semi-automated way, whereas for others it is hand-made by the authors. In the case of POWHEG [127, 128, 129], a framework is provided in which authors can provide plugins for various processes, some of which provide NLO QCD or even NNLO QCD predictions. An overview of the generators and the corresponding order in QCD used for signal and background event samples is presented in this section. Details related to the selection of background processes to consider can be found in Section 7.2.

The main task of the matrix element generator is to provide events sampled from the partonic cross section. In the majority of cases, the ME generator will also provide NLO QCD corrections to the partonic cross section, which fall into two categories: *virtual* corrections, which do not alter the ME-level final state, but do alter the amplitude via an additional parton connecting two existing partons in the LO diagram; and *real* corrections, where an additional parton is radiated, in which case the ME-level final state is altered. Example Feynman diagrams for virtual and real NLO correction terms are shown for the $q\bar{q} \to ZZ$ process in Fig. 5.1. In the case of real emission of a parton, either due to NLO corrections or due to multiple partons in the final state of the LO process, a *merging* scheme must be employed to prevent double-counting of parton (eventually, jet) radiation, since the parton shower may add partons to the ME-level final state in the process of hadronization. The merging of jets from matrix element calculations and parton shower descriptions is done using: the MLM [130] scheme in the case of LO samples with multiple partons in the final state produced

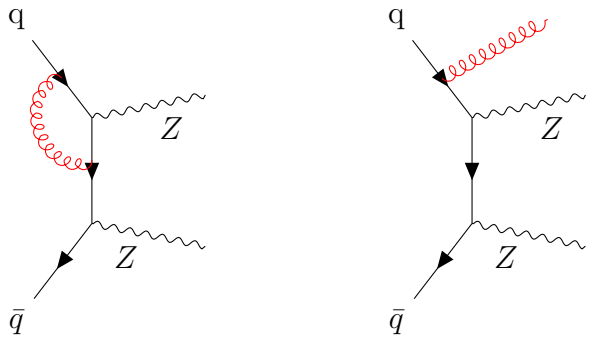

Figure 5.1: Example Feynman diagrams showing NLO QCD corrections to ZZ production, where an additional virtual gluon exchange (left diagram) or real gluon emission (right diagram) alters the amplitude of the ZZ process. In the case of virtual exchange, the strictly NLO contribution (i.e. one additional α_s factor) arises from the cross term with the LO diagram upon squaring the sum of amplitudes. The amplitude contributed by the sum of all NLO diagrams to the total ZZ cross section is subdominant to that provided by the LO diagram shown in Fig. 2.3.

with MadGraph5_aMC@NLO; the FxFx [131] scheme in the case of NLO samples generated by MadGraph5_aMC@NLO; and the MiNLO [132] scheme in the case of NLO samples generated by POWHEG. The choice of merging scheme is largely dictated by the supported options of the generator.

The POWHEG 2.0 generator is used to provide ME-level event samples for the following processes: diboson (VV, where V = W or Z) production via $q\bar{q}$ annihilation [133, 134]; ZH production via $q\bar{q}$ annihilation [135]; and ZH production via gluon fusion[2]. For all three processes, the ME calculation includes also NLO corrections. This choice reflects the state of the art in differential cross section predictions for these processes.

The Monte Carlo for FeMtobarn processes (MCFM) v7.01 [136] generator is used to produce event samples for the $gg \to$ WW and $gg \to$ ZZ processes, including also NLO QCD corrections. These so-called "box-diagram" processes can be viewed as NNLO QCD corrections to the respective LO processes, however they are found to

[2]No publication is available yet for POWHEG `User-Processes-V2/ggHZ`.

have a significant contribution due to the large gluon fraction of the proton PDF at the LHC energy scale. MCFM is chosen as the best prediction for these processes currently available. The $gg \to \mathrm{ZH}$ process is similarly enhanced relative to the naïve expectation for an NNLO correction. For this process, as mentioned previously, the POWHEG generator is used.

The reducible background samples, including Z+jets, Zγ, $\mathrm{t\bar{t}}$, $\mathrm{t\bar{t}V}$, and VVV, are generated at NLO in QCD with either MadGraph5_aMC@NLO v2.3.2 or POWHEG, depending on the process. The choice of generator for these reducible backgrounds is driven by the availability of computing resources and the requirements of analysis groups within the CMS collaboration that treat these processes as signals. Here we choose the sample which provides the lowest statistical uncertainty, provided it is known to model the data well in the relevant phase space.

Samples of DM particle production in the simplified model framework are generated using DmSimp [137, 138, 139] interfaced with MadGraph5_aMC@NLO v2.4.3. Samples are generated over a range of values for the masses m_{med} and m_{DM}. For the vector and axial vector models, the mediator couplings to the SM and DM fields are set to $g_q = 0.25$ and $g_{\mathrm{DM}} = 1$, respectively, and include full NLO corrections. For the scalar and pseudoscalar models, the couplings are set to $g_q = g_{\mathrm{DM}} = 1$, and the samples are generated only at leading order. This choice of couplings is recommended by the ATLAS/CMS dark matter forum [33] and by the LHC dark matter working group [35], and is motivated by the preference to explore new physics models with couplings initially set to unity, except that g_q is lowered in the vector and axial vector models such that pre-existing constraints from dijet searches (in the case where a DM mediator is produced via $q\bar{q}$ annihilation and decays back into a pair of quarks rather than DM particles) are evaded [140]. The DM and mediator mass ranges are chosen to be near the exclusion limits of current collider DM searches.

Events for the ADD scenario of large extra dimensions and for the unparticle model are generated at leading order using an EFT implementation in PYTHIA 8.205 [141, 142, 41]. In the ADD case, event samples are produced for $M_{\mathrm{D}} = 1$, 2, and 3 TeV, each with $n = 2$–7. The choice of M_{D} and n is driven by the expected sensitivity for these models. In order to ensure the validity of the EFT, the ADD signal is truncated following Eqn. 2.7. In general, this procedure has a larger effect for large values of n, for which the distribution of $\hat{s}$ is shifted towards higher values [142]. For the unparticle case, samples are generated for scaling dimensions $d_{\mathcal{U}}$ between 1.01 and 2.2, with the cutoff scale $\Lambda_{\mathcal{U}}$ set to 15 TeV and the coupling λ set to 1. Since both $\Lambda_{\mathcal{U}}$ and λ modify the cross sections of the signal prediction, but not its kinematic distributions [41], a simple rescaling of cross sections is performed to obtain signal predictions for alternative values of these parameters. No truncation is performed for the unparticle signal, to be consistent with the results of prior work.

5.4 Higher order corrections

The ME predictions for the ZZ and WZ processes are calculated at NLO in QCD, as discussed in the previous section. Higher-order calculations are known for both of these processes, some of which can be quite significant in certain differential distributions. To improve the description of these background processes, a reweighting procedure is used, as the higher-order cross sections are only computed differentially in a limited subset of the final state kinematic variables, preventing the calculation from being used as an event generator. These corrections include:

- a constant (approximately 10%) correction for the WZ cross section from NLO to NNLO in QCD calculations [143];
- a constant (approximately 3%) correction for the WZ cross section from LO to

NLO in electroweak (EW) calculations, considering also photon-quark initial states, according to Ref. [144];

- a $\Delta\phi(\mathrm{Z},\mathrm{Z})$-dependent correction, varying in magnitude up to 15%, to ZZ production cross section from NLO to next-to-next-to-leading order (NNLO) in QCD calculations [145];

- a p_{T}-dependent correction in the subleading Z boson p_{T}, varying in magnitude up to 20% at high $p_{\mathrm{T}}^{\mathrm{Z}}$, to the ZZ cross section from LO to NLO in EW calculations, following Refs. [146, 147, 144], which is the dominant correction in the signal region of this analysis.

In the case of WZ production, the corrections are applied to the total cross section. For ZZ production, events are given a weight w, formed from the ratios of differential cross sections k,

$$\begin{aligned} k_{\text{NNLO QCD}} &= \frac{\frac{d\sigma}{d\Delta\phi(\mathrm{Z},\mathrm{Z})}_{\text{NNLO QCD}}}{\frac{d\sigma}{d\Delta\phi(\mathrm{Z},\mathrm{Z})}_{\text{NLO QCD}}}, \\ k_{\text{NLO EW}} &= \frac{\frac{d\sigma}{dp_{\mathrm{T}}^{\mathrm{Z}}}_{\text{NLO EW}}}{\frac{d\sigma}{dp_{\mathrm{T}}^{\mathrm{Z}}}_{\text{LO}}}, \\ w &= k_{\text{NNLO QCD}} \cdot k_{\text{NLO EW}}, \end{aligned} \tag{5.3}$$

suitably binned in the available kinematic variables to capture the dominant effects of the higher-order prediction. The NLO EW correction is applied directly to the NLO QCD prediction, although the correction is with respect to the LO QCD prediction, under the assumption that the effects of the two corrections are factorisable.

5.5 Parton shower

To properly simulate the complete final state of the collision, the interactions of the underlying event, the evolution of ME partons down to the hadronization energy scale,

and the subsequent hadronization of partonic remnants of the collision—including both those of the underlying event as well as any final state partons of the hard scatter—must be simulated. In particular, if a color current (a conserved current of the $SU(3)$ symmetry) is carried through the hard scatter, then the color connection between the hard scatter final state and the proton remnants must also be considered. The evolution of ME partons to low energy scales is accomplished via *splitting functions* [148], a result of the DGLAP equations that describes the probability of a parton to evolve from a high energy scale to a lower one without radiating an additional parton above some minimum energy cutoff, where the cutoff is necessary as the splitting function is generally infrared and collinear divergent, i.e. divergent in the transverse momentum and opening angle of the radiated parton relative to the initial parton [149]. After all partons are near the hadronization energy scale (about 10 GeV), the modeling of the underlying event and hadronization becomes largely phenomenological. Three popular parton shower programs exist, each with different phenomenological treatments of hadronization: PYTHIA [141], HERWIG [150], and SHERPA [151]. Each model contains many free parameters that are adjusted, or *tuned*, to best match experimental data. In this analysis, PYTHIA versions 8.205 or higher is used with all event samples for parton showering, hadronization, and the underlying event simulation, using the CUETP8M1 [152] tune.

5.6 Detector response

The final step of event simulation is tasked with simulating the detector response to the final state particles of the proton-proton collision, which is handled by the GEANT4 [153, 154, 155] program, in combination with custom digitization simulation software. The event description at the exit of the parton shower simulation contains

several thousand particles emanating from the interaction vertex, generally consisting of electrons, muons, photons, light stable mesons such as pions and kaons, and stable baryons such as protons or neutrons. Some particles are unstable with lifetimes of order $c\tau \sim$ m, whose decays are handled by GEANT. The particles are propagated through a simulation of detector material, where all known matter interactions are considered, weighted by their respective probabilities using the Monte Carlo method, and energy deposits in each material volume are recorded. The energy deposits in sensitive volumes are converted into digital values as would be read out by the detector electronics, using models developed and validated with both test beam [156] and collision [157, 158] data. These models capture the effects of charge or light collection efficiency, amplifier noise, pulse shaping circuitry, and many other aspects of the detector response. The readout circuitry, in particular that which generates the trigger primitives and selective readout (zero-suppression), is simulated as well.

During the digitization step, events are superimposed with *minimum bias* events—generic low momentum transfer inelastic scattering events that approximate the average proton-proton collision—to simulate the effect of pileup. The minimum bias events are simulated using PYTHIA. Up to 70 minimum bias events may be superimposed with the hard scatter event, where the number of interactions is drawn from a distribution selected to approximate that of the data collection period for which the simulated events are to be used with. At the analysis level, once the dataset is finalized, the pileup distribution of events in simulation is reweighted to exactly match that of the data. The average number of pileup events per bunch crossing is approximately 23 in the data sample analyzed here, and the distribution is shown in Fig. 5.2.

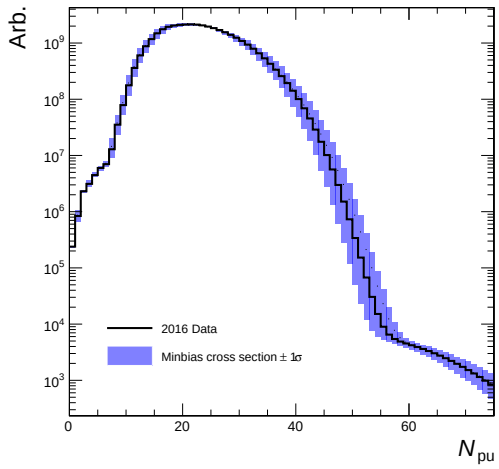

Figure 5.2: Distribution of pileup observed in data collected in 2016, with a mean $N_{\text{pu}} = 23$. N_{pu} is constructed from the sum of per-bunch instantaneous luminosity distributions, recorded per lumisection, scaled by a minimum bias cross section of $69.2 \pm 3.5\,\text{mb}$ [102].

Chapter 6

Event reconstruction

6.1 Introduction

The event reconstruction algorithms employed by CMS have been in constant development over many decades, in response to the increasingly difficult reconstruction challenges posed by successive generations of particle detector experiments. The algorithms have increased in complexity in step with the experiments themselves. A common theme in algorithm development is iterative improvements: initial simplified algorithms feed more refined and more computationally complex algorithms, such that the combinatorial phase space of 70 million channels of CMS tracking detector and calorimeter readout is reduced in stages. Each step increases the purity as much as possible, within acceptable efficiency losses, of the main physics objects in an event. CMS has successfully employed the particle flow (PF) technique [159] as a core feature of reconstruction. The goal of PF is to reconstruct all particles produced in the event, divided into five classes of visible particles present in the detector: muons, electrons, charged hadrons, neutral hadrons, and photons. A schematic of how each class of particles leave traces in the detector is shown in Fig. 6.1. PF reconstruction

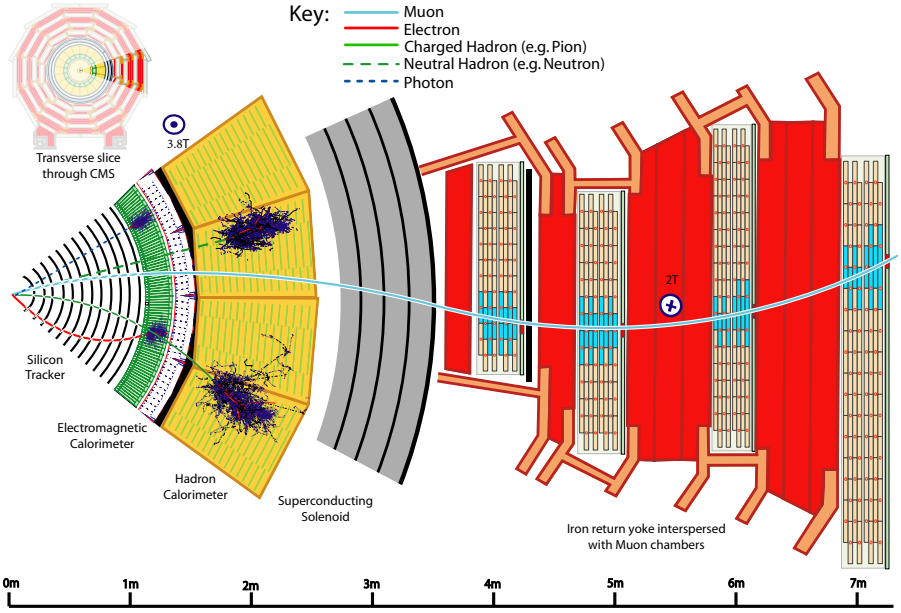

Figure 6.1: A sketch of the specific particle interactions in a transverse slice of the CMS detector, from the beam interaction region to the muon detector. The muon and the charged pion are positively charged, and the electron is negatively charged. Reproduced from [159], Fig. 1.

starts with local algorithms that form tracks and calorimeter clusters, then a linking stage combines these into PF objects. From the visible particles reconstructed by PF, the main physics objects used in this analysis are then formed and identified: isolated muons, isolated electrons, isolated photons, jets, isolated tau leptons, and the missing transverse momentum of the event, $p_{\mathrm{T}}^{\mathrm{miss}}$.

6.2 Tracks

6.2.1 General tracks

Track reconstruction begins with the local reconstruction of clusters in the pixels and strip trackers (Section 4.5), where the zero-suppressed signals collected by the

detectors are clustered into *hits*. The estimated positions and uncertainties of the hits, in the plane of the sensor, is computed [160]. The efficiency for a hit in a given layer to be reconstructed and used in a track fit, excluding damaged sections of the detector, is over 99.5% for both strip and pixel hits. In 2016, about 98% of the pixel detector channels and about 96% of the strip tracker channels were active and functional. Hit positions in this local coordinate system are then translated into the global CMS coordinate system via a geometry specification that is informed by *in-situ* alignment measurements that correct for deviations from the ideal positioning of the detector elements as well as surface deformations due to mechanical stresses. The statistical accuracy of the alignment procedure is found to be less than $10\,\mu$m [161].

Prior to track reconstruction, the position of the beam spot, i.e. the center of the luminous region in which the proton bunches collide, must be established. The beam spot is measured using only pixel hits, fed to a simplified tracking algorithm which generates and clusters tracks to determine primary vertex positions, which are then fit to a 3D Gaussian beam profile to extract the beam spot position over an average of many events [160]. The beam spot is used to constrain track seeding for an initial subset of the tracking iterations.

Track reconstruction is performed in ten iterations, beginning with the reconstruction of high-quality, high-p_{T} tracks originating from the beam spot, and progressively relaxing quality criteria and displacement (from the beam spot) thresholds in each iteration. Hits used in one iteration are removed from the collection available to the next iteration to reduce computational complexity as well as false track reconstruction probability. Each iteration follows four steps. First, a set of *seeds*, or collections of hits in two to three detector layers, is defined, which provides the initial estimate of the track trajectory parameters and their uncertainties. Next, compatible hits are searched for by extrapolating the particle trajectory (outward or inward) into the

next layers using the Combinatorial Track Finder (CTF) method, which is a type of Kalman filter [160, 162, 163]. Following the building of CTF tracks, a track fit is performed to provide the best possible estimate of the track parameters. Finally, track selection discards track that fail certain quality criteria. The first three track-finding iterations are seeded by triplets of pixel hits, capturing most prompt high-p_{T} tracks, as well as mildly displaced tracks from b decays. The fourth step uses only pairs of pixel hits consistent with the beam spot to recover the remaining high-p_{T} tracks. Iterations five through seven target progressively further displaced tracks, with origin vertices as far as 60 cm from the beam spot. The eighth iteration attempts to resolve merged tracks in the core of high-p_{T} jets by splitting hit clusters as necessary. The last two iterations are seeded by information from the muon detectors, with outward- and inward-propagating track finding steps. The efficiency of track finding and the track misreconstruction rate for this iterative method is shown in Fig. 6.2.

6.2.2 Electron tracks

For electrons, unlike for muons or charged pions, there is a high probability to emit significant bremsstrahlung radiation as the electron traverses the significant (up to 1.8 radiation lengths) tracker material before entering the electromagnetic calorimeter. As a consequence, electron tracks can have sharp kinks, at which point the CTF tracking procedure will fail to find the next hit. Therefore, a dedicated electron tracking procedure, known as the Gaussian-Sum Filter (GSF) method [164] is employed. Seeding of electron tracks is performed using two sets of inputs: calorimeter superclusters (Section 6.4) and general tracks. In the first case, GSF tracking is seeded via isolated superclusters with $E_{\mathrm{T}} > 4\,\mathrm{GeV}$ and a hadronic to electromagnetic energy ratio below 0.15, where the initial momentum is assigned by back-propagating the supercluster position under both electron and positron hypotheses. For prompt iso-

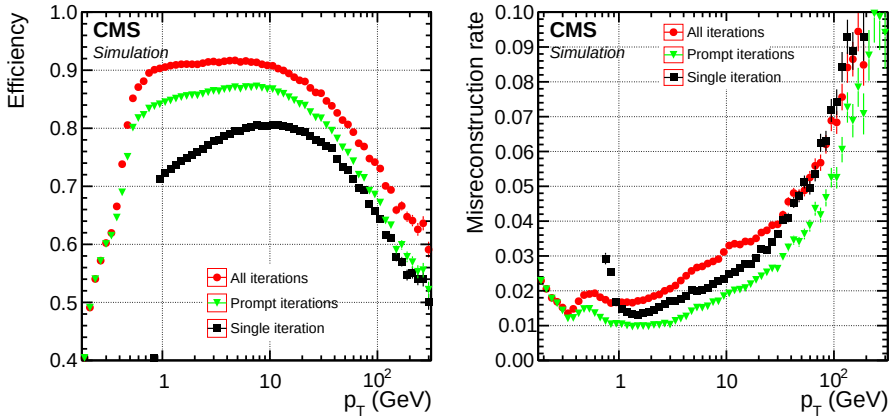

Figure 6.2: Efficiency (left) and misreconstruction rate (right) of the global combinatorial track finder (black squares); and of the iterative tracking method (green triangles: prompt iterations based on seeds with at least one hit in the pixel detector; red circles: all iterations, including those with displaced seeds), as a function of the track p_{T}, for charged hadrons in multijet events without pileup interactions. Only tracks with $|\eta| < 2.5$ are considered in the efficiency and misreconstruction rate determination. The efficiency is displayed for tracks originating from within 3.5 cm of the beam axis and ± 30 cm of the nominal center of CMS along the beam axis. Reproduced from [159], Fig. 4.

lated electrons from Z boson decays, this seeding method alone has an efficiency of about 92% [165]. For non-isolated or low-p_{T} electrons, as well as those in the barrel-endcap transition region, the supercluster requirements become too restrictive, so an additional GSF reconstruction sequence is seeded by all general tracks over 2 GeV. The subset of electrons with minimal bremstrahlung will already be well-measured by the CTF tracking and can later be matched to calorimeter clusters to discriminate against charged hadrons (Section 6.5), so to recover those with large radiation, the subset of CTF tracks with a low number of hits and high fit χ^2 are used to seed a GSF tracking iteration. Electron candidates from both seeding methods are then merged and used to seed a more robust and computationally expensive GSF tracking iteration to determine the final electron tracks. The electron track reconstruction efficiency is above 95% for electron superclusters with $E_T > 20$ GeV and $|\eta| < 2$ [165].

6.2.3 Muon tracks

Muon tracks fall into three categories: stand-alone, tracker, and global muons. Stand-alone muon tracks are built using only information from the muon detectors, where hits are first locally reconstructed within DT and CSC detector elements, then formed into track segments, which then seed a Kalman filter track finding algorithm to combine all DT, CSC, and RPC hits compatible with the track [166, 167]. The efficiency for reconstructing local segments and hits is on average about 95% per muon station in the CSC and DT systems, and the detection efficiency of the RPC layers is typically between 95-98% [168]. Tracker muons include any general tracks with $p_{\mathrm{T}} > 0.5\,\mathrm{GeV}$ and total momentum $p > 2.5\,\mathrm{GeV}$ that, when extrapolated to the muon system, can be matched to at least one muon segment within $min(3\,cm, 4\,\sigma)$ in the best-measured local coordinate in the plane transverse to the beam axis. Global muons are formed by extrapolating general tracks and stand-alone muon tracks to a common surface (namely, the tracker outer surface) and matching tracks which are compatible. The hits from the matched tracks are then re-fit, which provides a significant improvement in the momentum resolution for muons with $p_{\mathrm{T}} > 200\,\mathrm{GeV}$ [169]. Since global muon reconstruction requires hits in multiple muon detector layers, the global muon reconstruction efficiency reaches a plateau efficiency of 99% at a muon p_{T} near $10\,\mathrm{GeV}$, whereas tracker muons are reconstructed with 99% efficiency above $1\,\mathrm{GeV}$ [169].

6.3 Primary vertices

Primary vertices are the vertices from which tracks associated with initial proton-proton collisions emanate, to be contrasted with secondary vertices, from which tracks associated with particle decays or interactions with detector material emanate. The goal of primary vertex reconstruction is to measure the position of all proton-proton

collisions in the event, both from the hard scatter as well as from pileup collisions. Primary vertices are reconstructed in three steps: a subset of the reconstructed track collection is selected based on compatibility with the beam spot and track quality criteria, without any restriction on track p_{T}; the tracks are clustered according to the z coordinate and the point of closest approach to the center of the beam spot, using the deterministic annealing [170] algorithm; and the position of each vertex is then found by a maximum likelihood fit to the associated tracks belonging to a given cluster [160].

The selected primary vertex of the event is primary vertex associated with the specific proton-proton collision that produced a hard scatter event, i.e. not a pileup collision. In run 1, the selected vertex was that with the largest sum of squared transverse momenta of associated tracks, $\sum p_{\mathrm{T}}^2$. For some classes of signal models, such as those where substantial transverse energy is carried by neutral or invisible particles, this selection was found to be non-optimal at high pileup. To improve the selection efficiency in run 2, the selected vertex is that with the largest value of summed physics-object p_{T}^2, where in this context the physics objects are: track-jets, clustered using the AK4 jet-finding algorithm (see Section 6.6.4) with the tracks assigned to the vertex as inputs; and the associated missing transverse momentum, taken as the negative vector sum of the p_{T} of the track-jets associated with the vertex. This procedure improves the selected primary vertex efficiency in processes such as $\mathrm{H} \to \gamma\gamma$ and $\mathrm{ZH} \to \nu\nu b\bar{b}$ by about 5%, while for the signals considered here, the efficiency is not altered significantly from its already high ($> 98\%$) value as this final state contains two high-p_{T} tracks from the $\mathrm{Z} \to \ell\ell$ decay.

6.4 Calorimeter clusters

Calorimeter clustering is designed to accomplish several tasks. Photons and neutral hadrons are only detectable in the calorimeters, so the clustering must accurately measure the energy and position of these neutral particles. Electron tracking is seeded by superclusters, which are contiguous groups of ECAL clusters designed to capture bremsstrahlung photons emitted by the electrons. Clusters are also used to improve the energy measurement of high-p_{T} charged hadrons in the case where the track associated to the hadron has a poor momentum measurement. Calorimeter clusters are formed separately in each subdetector: ECAL barrel and endcaps, HCAL barrel and endcaps, and the ECAL preshower. For HF, no clustering is performed, and each detector element becomes itself a cluster.

The initial clustering is topological: a detector element with energy deposit above a threshold E_{seed} defines a seed element, and the cluster is grown around the seed by successively including adjacent elements to existing cluster elements (either the four sharing a side or the eight sharing a corner) which have energy deposits above a threshold E_{gather}. In the case of the ECAL endcaps, since the noise level increases strongly with $|\eta|$, an additional transverse energy threshold on the seed element, $E_{T,\mathrm{seed}}$, is imposed. The values of these clustering parameters are summarized for each detector in Table 6.1. After the topological clustering, an iterative algorithm refines the position and energy of clusters. The first step assigns a Gaussian energy distribution to each seed, with initial energy and position measurements given by the seed element, and the Gaussian width a fixed parameter chosen according to the detector in which the cluster resides. The clustering continues by performing an analytical maximum-likelihood fit of each cluster's energy to the energy deposit distributions, and then updating the position measurements and energy fractions of

Clustering parameter	ECAL		HCAL		Preshower
	barrel	endcaps	barrel	endcaps	
E_{seed} (MeV)	230	600	800	1100	0.12
$E_{T,\text{seed}}$ (MeV)		150			
Gathering adjacent cells	8	8	4	4	8
E_{gather} (MeV)	80	300	800	800	0.06
Gaussian width (cm)	1.5	1.5	10	10	0.2

Table 6.1: Clustering parameters for the ECAL, HCAL, and the preshower detectors. All values result from optimizations based on simulation of single photons, π^0, K^0_L, and jets. Reproduced with modified labeling from [159], Table 2.

each element per cluster before the next iteration. In this way, the energy deposited by two nearly-overlapping clusters is divided between them according to the expectation of the Gaussian model.

6.5 Particle-flow linking

As shown in Fig. 6.1, a given particle is expected to give rise to several PF elements in the various CMS subdetectors. With the track and cluster PF elements assembled, PF reconstruction proceeds with a *link algorithm* that connects these elements from different subdetectors to form the PF candidates. Each track is propagated into the position of expected shower maximum in ECAL, and linked to an ECAL cluster if the track position is within the cluster boundary plus one crystal, in the η-ϕ plane in the barrel, and in the x-y plane in the endcap. The tracks are further propagated one hadronic interaction length deep into the HCAL, and matched in the same way to HCAL PF clusters. If multiple track-cluster links are found, only the closest link is kept, with the distance between the track and cluster barycenter as the metric. Preshower, ECAL, and HCAL clusters are linked if the cluster in the detector with finer granularity (in order: PS, ECAL, HCAL) overlaps projectively with the other cluster. As in the track-calorimeter matching, only the closest link is kept. To

collect photons emitted by electron bremsstrahlung radiation, tangents to electron GSF tracks at each tracker layer are propagated to the ECAL, and links are formed between the GSF track and clusters intercepted by these tangents. Combinations of tracks that are found to intersect and form a common displaced vertex due to nuclear interactions in the tracker material are also linked together. Finally, tracks in the inner tracker and muon systems are linked. The end result is several PF *blocks* of objects linked together by the above algorithm. Due to the fine granularity of the CMS detector, typical PF blocks contain only a handful of elements originating from one or a few particles [159]. After the linking step, each PF block is then analyzed individually to determine the particle content in a specific order, chosen so that objects with intrinsically higher purity are removed from consideration at the beginning of the sequence.

Isolated muons are identified by the presence of linked inner and outer tracks (as well as a small calorimeter deposit due to the minimally-ionizing nature of muons) and by an isolation criterion requiring the sum of track p_{T} and calorimeter E_{T} in a $\Delta R < 0.3$ region to be less than 10% of the muon p_{T}, where ΔR is defined:

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}.$$

Non-isolated muons (e.g. from semileptonic B meson decays in b jets) are identified by more strict criteria, similar to the *tight* identification requirements described in Section 6.6.1. PF elements associated with the resulting muons are then masked from the PF block.

Next, electrons and photons are considered simultaneously, due to the similarities in their appearance: often photons will convert (interact with the tracker material, producing an e^+e^- pair), and bremsstrahlung photons from electrons may convert, leaving many e/γ objects from one originating hard photon or electron. In a given PF

block, an electron candidate is seeded by a GSF track, provided the corresponding ECAL cluster is not linked to more than 3 tracks, and a photon candidate is seeded by an ECAL supercluster with $E_{\mathrm{T}} > 10\,\mathrm{GeV}$ and no link to a GSF track. The total energy of the ECAL clusters linked to an e/γ object is corrected for energy lost due to inefficient clustering and linking [165, 159]. This corrected energy is assigned to final photon candidates. The final electron candidate energies are obtained from a multivariate regression using a combination of the corrected ECAL energy and the GSF track momentum as inputs, along with other track and cluster variables. The photon position is taken from the supercluster, and the electron position taken from the GSF track. Electron candidates are then filtered by a 14-variable multivariate discriminator, trained separately for barrel and endcaps and for isolated and non-isolated electrons [159]. Photon candidates are filtered by isolation criteria and the ratio of HCAL and ECAL energies. All tracks and clusters in the PF block used to reconstruct these e/γ objects are then masked.

After the removal of $\mu/e/\gamma$ candidates, the remaining particles to be identified in a given PF block are hadrons and non-isolated photons from jet hadronization. Photons and neutral hadrons (h^0) are formed from any ECAL and HCAL clusters not linked to any track, where inside the tracker region all ECAL clusters are assigned to γ candidates and all HCAL clusters to h^0 candidates. This choice is motivated by the observation that 25% of hadronic jet energy is carried by photons (primarily from π^0 decays) while neutral hadrons only deposit 3% of the jet energy in the ECAL. Outside the tracker acceptance ($|\eta| > 2.5$), charged and neutral hadrons become indistinguishable, therefore linked ECAL and HCAL clusters form neutral hadron candidates and unlinked ECAL clusters form photon candidates. The remaining HCAL clusters are linked to one or several tracks to form charged hadron ($h^{\pm}$) candidates. Calibrations are applied to the calorimeter clusters as a function of E, η, and the candidate hypoth-

esis: γ, $h^{\pm}$, or h^{0}. Finally, HF clusters directly form *HF photon* or *HF hadron* particle candidates, depending on the ratio of long to short fiber energies in the cluster.

After the complete PF particle collection is established, an event post-processing step checks for spurious $p_{\mathrm{T}}^{\mathrm{miss}}$ (Section 6.6.6), created for example by: reconstruction of a coincident cosmic ray muon in the event; or mis-reconstruction of a charged hadron as a muon and neutral hadron in the case that the $h^{\pm}$ is energetic enough for its shower to produce segments in the muon systems (referred to as *punch-through*), which causes a double-counting of the $h^{\pm}$ energy. Particles likely to cause such issues are masked from the $p_{\mathrm{T}}^{\mathrm{miss}}$ calculation, and if the resulting $p_{\mathrm{T}}^{\mathrm{miss}}$ value is reduced by a factor of at least 0.5, the particles are removed from the output collection or the PF block is re-interpreted as necessary.

6.6 Physics objects

6.6.1 Muons

Muons selected as primary objects in this analysis are reconstructed by the particle flow algorithm, and are subjected to additional identification criteria: the muon track is required to be globally reconstructed, with a track fit $\chi^2/\mathrm{d.o.f.} < 10$; the track fit must use at least one muon chamber hit, and there must be compatible segments (that may not have been used in the final fit) in at least two muon stations; the track position of closest approach to its associated vertex must be less than 0.2 mm in the radial coordinate and less than 1 mm in the z coordinate; the track must include at least one pixel hit, and the inner track must use hits in at least 6 layers for p_{T} measurement. These requirements primarily suppress non-prompt muons from B meson decays, charged hadron punch-through, and cosmic muons.

In this analysis, muons are expected to be well-isolated from hadronic activity,

as they are produced in the decay of a Z boson. To suppress backgrounds where muons are produced in semileptonic decays of heavy mesons in jets, a strict isolation requirement is imposed, using the ratio of PF particle p_{T} sums in a radius $\Delta R < 0.4$, with a correction for the pileup particle contribution given by the formula,

$$I^{\mu} = \sum_{\Delta R<0.4} p_{\mathrm{T}}^{h^{\pm}} + \max\left(\sum_{\Delta R<0.4} p_{\mathrm{T}}^{h^{0}} + \sum_{\Delta R<0.4} p_{\mathrm{T}}^{\gamma} - \frac{1}{2}\sum_{\Delta R<0.4} p_{\mathrm{T}}^{h^{\pm,PU}}, 0\right), \tag{6.1}$$

where $h^{\pm,PU}$ are PF charged hadrons not associated with the selected primary vertex of the event. The factor of one half encodes the expected charged to neutral energy ratio in a typical pileup jet, thereby subtracting the expected pileup contribution from the neutral energy sums. In this analysis, muons are required to have $I^{\mu}/p_{\mathrm{T}}^{\mu} < 0.15$.

For the purpose of rejecting events with additional muons beyond those expected in a given channel (see Section 7.4), a relaxed selection is used to veto events with additional muons, where the identification requirement is that the PF muon is reconstructed as a global or tracker muon (i.e. not a stand-alone muon), and the isolation value is relaxed to $I^{\mu}/p_{\mathrm{T}}^{\mu} < 0.25$.

The sign of the muon charge is determined by the direction of curvature of the muon track in the magnetic field, and the muon charge mis-assignment probability is found to be negligible in both data and simulation [171].

The efficiency of muon reconstruction, identification, and isolation is measured in data with the "Tag and Probe" technique [169], where a selection targeting events with $\mathrm{Z} \to \mu\mu$ decays is made, but only considering one of the muon daughters (the *tag*) so that the properties of the other muon daughter (the *probe*) can be evaluated in an unbiased way. To do so, the tag muon is selected with requirements very similar to the above identification, and additionally required to be a triggered muon, i.e. it has sufficient energy and quality to pass a single muon trigger and cause the event to be collected, as opposed to an event collected by a double muon trigger

path, where both the tag and probe muon affect the data collection efficiency. Then, general tracks in the event are each paired with the tag muon, and the invariant mass is computed and placed in a histogram. Then, if the track is matched to a fully identified muon, a "passing probe" histogram is filled. These histograms are then fit to a Z boson line shape function plus a generic function to model the combinatorial background from probe tracks that are not associated with the Z boson resonance, to extract the number of probe muons from true Z boson decays passing and failing the identification requirements. Thus, the efficiency of true muons to pass the identification requirements is measured, with respect to the efficiency to reconstruct a general track for true muons, which is measured independently to be better than 99% for isolated muons with $1 < p_{\mathrm{T}} < 100\,\mathrm{GeV}$ within the entire tracker acceptance [160]. The same efficiency measurement procedure is also carried out in simulation, and the simulation samples are re-weighted by the ratio of data and simulation efficiencies per selected muon in the event. The typical magnitude of these corrections is 1-3% for the dataset used in this analysis [167]. Fig. 6.3 shows the efficiency of muon identification and isolation measured in data and simulation for the dataset used in this analysis.

6.6.2 Electrons

Electrons selected as primary objects in this analysis are reconstructed by the particle flow algorithm. However, they are not subjected to the multivariate discriminator filtering of the PF event description but rather to identification criteria based on observables related to the GSF track, the clustered calorimeter deposits, and the track-cluster compatibility. The values of the various thresholds for the parameters described below are summarized in Table 6.2. The GSF track and ECAL supercluster positions are checked for compatibility via selections on $|\Delta\eta_{\mathrm{seed}}|$ and $|\Delta\phi_{\mathrm{in}}|$, which are,

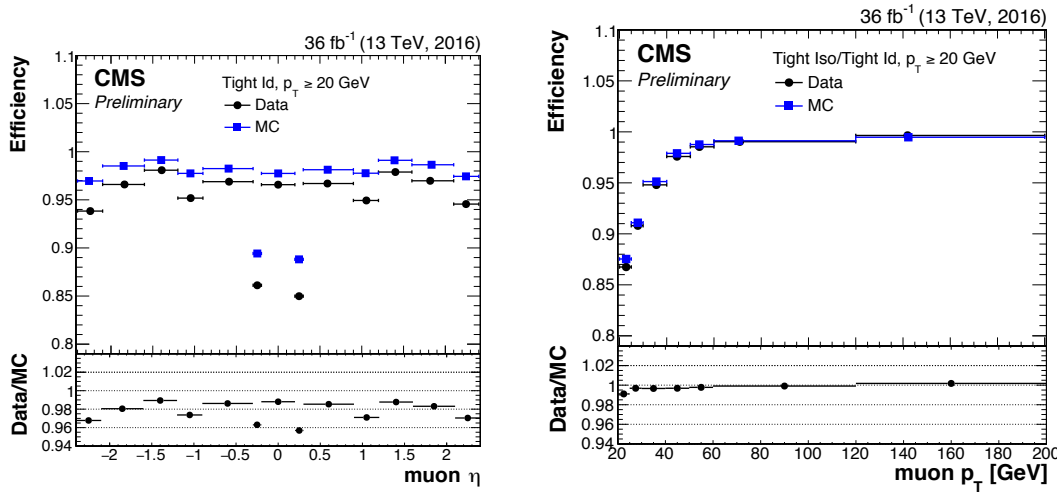

Figure 6.3: The efficiency of the signal muon selection as measured in data and simulation for this analysis. Left: muon identification efficiency with respect to general track reconstruction efficiency, as a function of muon η, for muons with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$. Right: isolation efficiency for muons from Z decays with respect to reconstructed and identified muons in the $|\eta| < 2.4$ acceptance of the CMS muon systems. Reproduced from [172].

respectively, the difference between the GSF track η at the vertex and the η position of the seed ECAL PF cluster that forms the supercluster, and the difference between the ϕ value of the GSF track at the origin and the ϕ value of the supercluster. The track and supercluster must also match well in energy, via a selection on $|1/E_{\mathrm{SC}} - 1/p_{\mathrm{GSF}}|$, where E_{SC} is the supercluster energy and p_{GSF} is the total momentum of the GSF track. To reject photon conversions, the track must have at most one missing hit in the active regions of the pixel detector, and the track must not be associated with any reconstructed conversion vertices. Hadrons are rejected by a selection on the ratio of hadronic to electromagnetic energy deposits H/E that is a function of the energy of the supercluster E_{SC} and the median energy density in the event ρ, as computed by averaging the PF particle energy distribution over the geometrical coverage of the detector. The shower shape of the supercluster must be compatible with that of an electron, via a selection on $\sigma_{i\eta i\eta}$, defined in a 5×5 crystal region centered at the seed

crystal of the seed cluster of the supercluster, as

$$\sigma_{i\eta i\eta} = \sqrt{\frac{\sum_{i\in 5\times 5} w_i (i\eta_i - \bar{i\eta}_{5\times 5})^2}{\sum_{i\in 5\times 5} w_i}}, \qquad w_i = 4.7 + \ln\left(\frac{E_i}{E_{5\times 5}}\right), \tag{6.2}$$

where E_i is the energy of the crystal, and $i\eta$ is the integral value of the η position of the crystal for barrel crystals, and is defined as $i\eta = \sqrt{ix^2 + iy^2}$ in the endcap, where the crystal arrangement follows an x-y grid rather than an η-ϕ grid. Finally, the electron is required to be isolated, where the isolation is defined as

$$I^e = \sum_{\Delta R<0.3} p_{\mathrm{T}}^{h^\pm} + \max\left(\sum_{\Delta R<0.3} p_{\mathrm{T}}^{h^0} + \sum_{\Delta R<0.3} p_{\mathrm{T}}^{\gamma} - \rho \cdot A_{eff}(|\eta_{\mathrm{SC}}|), 0\right), \tag{6.3}$$

where $A_{eff}(|\eta_{\mathrm{SC}}|)$ is the effective area of the electron footprint, which, when multiplied by the median energy density ρ, provides an estimate of the local pileup contribution to the neutral energy sum. The specific values of the above selections are optimized for signal efficiency and background rejection using simulated samples, separately for barrel and endcap electrons, and several working points are defined. In this analysis, the "medium" working point is used to select signal electrons, designed to target a true electron identification efficiency of approximately 80%, and the "veto" working point is used to reject events with additional electrons beyond those expected in a given channel, targeting a true electron identification efficiency of 95%.

Due to the high probability of conversions and bremsstrahlung for electrons, the sign of the electron charge is determined by a majority vote of three charge estimators: the direction of curvature of the GSF track in the magnetic field, the direction of curvature of the associated CTF track in the case that a hit is shared between the two tracks, and the sign of the difference in ϕ between the vector joining the beam spot to the supercluster position and the vector joining the beam spot to the first hit of the electron GSF track. With this majority method, the electron charge mis-assignment

Electron ID parameter	Working point in barrel (endcaps)			
	Signal selection		Veto selection	
$\lvert\Delta\eta_{\text{seed}}\rvert$	0.00311	(0.00609)	0.00749	(0.00895)
$\Delta\phi_{\text{in}}$	0.103	(0.045)	0.228	(0.213)
$\lvert 1/E_{\text{SC}} - 1/p_{\text{GSF}}\rvert$	0.134	(0.13)	0.299	(0.15)
Missing inner track hits	1		1	(3)
Conversion veto	Yes		Yes	
H/E	0.253	(0.0878)	0.356	(0.211)
$\sigma_{i\eta i\eta}$	0.00998	(0.0298)	0.0115	(0.037)
I^e/p^e_{T}	0.0695	(0.0821)	0.175	(0.159)

Table 6.2: Threshold values for electron identification and isolation working points used in this analysis, for the parameter definitions as described in Section 6.6.2. The threshold value used in the endcap regions is shown in parenthesis, where it differs from that of the barrel region.

probability is about a factor of 2 better than that of the GSF track curvature alone, and is predicted in simulation to be 1.5% [165].

The efficiency of electron reconstruction, identification, and isolation is measured in data with the same technique as for muons, exploiting the di-electron decay mode of the Z boson. The efficiencies are found to be in good agreement between data and simulation, with typical correction factors of order 1-5%, and compatible with unity within the uncertainties of measurement [165]. Fig. 6.4 shows the efficiency of electron GSF track reconstruction and identification measured in data and simulation for the dataset used in this analysis.

6.6.3 Photons

Photons are not selected as primary objects in this analysis, however they do enter into the $p_{\text{T}}^{\text{miss}}$ calculation as PF photon objects, both isolated and non-isolated. The efficiency of photon reconstruction is measured in data using $\text{Z} \to ee$ events where the electron track is ignored and hence is reconstructed as a photon, as well as in $\text{Z} \to \mu\mu\gamma$ events where a high-energy photon is emitted from the final state muon,

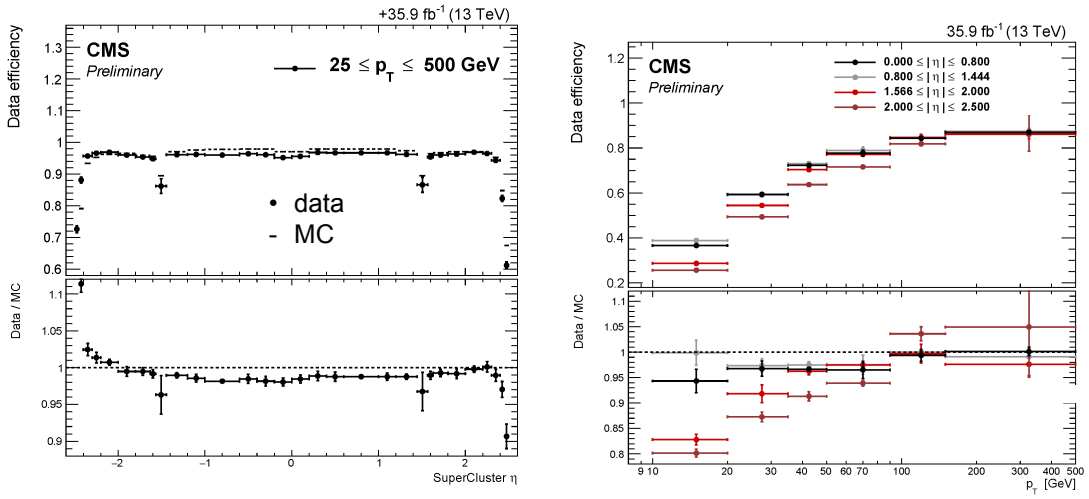

Figure 6.4: The efficiency of the signal electron selection as measured in data and simulation for this analysis. Left: electron GSF track reconstruction efficiency with respect to reconstructed superclusters, as a function of supercluster η, for electrons with $p_{\mathrm{T}} > 25\,\mathrm{GeV}$. Right: electron identification efficiency with respect to reconstructed electrons, as a function of electron p_{T} for various $|\eta|$ regions, for an electron identification working point more restrictive than the one used in this analysis. (No public plot is available for the working point used in this analysis.) Reproduced from [173].

and is found to be well-modeled by simulation [174].

6.6.4 Jets

Jet objects are composite objects, made up of clusters of PF candidates. The ideal clustering algorithm is designed with the intrinsic properties of parton showers (as discussed in Section 5.5) in mind—in particular, it is desirable to be insensitive to the infrared and collinear divergences present in the theoretical calculations of differential cross sections for processes of interest. Several such algorithms have been identified (a thorough review can be found in Ref. [175]), and the anti-k_t algorithm [176] with distance parameter $R = 0.4$ (hereafter, AK4) has emerged as the preferred choice in recent years, due mainly to the fact that the jet boundaries are resilient to soft radiation, and the consequences that follow, e.g. it produces jets with roughly conical

catchment areas, which simplifies experimental calibration procedures. The choice of distance parameter $R = 0.4$ is found to be a good compromise between dijet mass resolution (preferring larger R for dijet resonances in the TeV range [175]) and the amount of pileup particles captured (preferring smaller R.) The AK4 algorithm can be described as follows: starting with all PF particles in the event, generate a collection of *pseudojets* from the momenta of each particle and compute the distance parameters d_{ij} between each pair of particles and d_{iB} between each particle and the beam as follows,

$$\begin{aligned} d_{ij} &= \min(k_{t,i}^{-2}, k_{t,j}^{-2}) \frac{\Delta_{ij}^2}{R^2}, \\ d_{iB} &= k_{t,i}^{-2}, \end{aligned} \tag{6.4}$$

where $\Delta_{ij}^2 = (y_j - y_i)^2 + (\phi_j - \phi_i)^2$, and $k_{t,i}$, y_i, and ϕ_i are respectively the transverse momentum, rapidity, and azimuthal angle of pseudojet i. The distance parameters are sorted, and if the smallest parameter is a d_{ij}, the pseudojets i and j are merged into a pseudojet by summing their momenta, or if the smallest parameter is a d_{iB}, pseudojet i is removed from consideration. Then, the procedure is repeated until no pseudojet entities are left. The resulting particles that did not get merged are referred to as *unclustered* particles, and the resulting pseudojets become the final jet collection. In the implementation used for this analysis, the AK4 algorithm is slightly modified, where any charged hadron that is not associated with the chosen primary vertex of the event ($h^{\pm,PU}$) is discarded prior to the clustering, referred to as the charged-hadron subtraction (CHS) procedure [177]. Despite the fact that neutral particles from pileup are still clustered, AK4 CHS jets are found to be more resilient to the effects of pileup contributions [178]. The neutral pileup contribution to the jet energy is corrected for by subtracting a local pileup contribution estimate based on the jet area and the median energy density ρ [179]. Additional energy corrections are applied to the jets used in this analysis to account for detector response and inefficiency, as

measured in simulation by comparing jets formed by clustering stable (decay length $c\tau > 1\,\text{cm}$) visible final state particles at generator level to the reconstructed jets; and residual data to simulation discrepancies, as measured in data via techniques that exploit the expected energy balance of dijet events [179]. The calibrated jet energy scale uncertainty is better than 2.5% for jets with $p_{\text{T}} > 30\,\text{GeV}$ and $|\eta| < 2.5$ for the dataset used in this analysis [180].

Jets used in this analysis are filtered to remove detector noise, via selections on the multiplicity and energy fraction of the various PF candidate constituents, as true jets are expected to contain multiple candidate species with no single species contributing the entirety of the jet energy. The "loose" jet identification [178] used in this analysis has better than 99% efficiency for true jets up to $|\eta| < 3$, and better than 95% efficiency for forward jets, as measured in data via a tag and probe technique in dijet events [178].

Jets originating from the hadronization of bottom quarks (b-jets) have a distinguishing characteristic: as many B mesons have relatively long lifetimes (e.g. $c\tau$ for $B^{\pm}$ is 491 μm [9]), the charged hadrons from the eventual decay of B mesons in a b-jet will originate from a slightly displaced vertex. A dedicated b-tagging algorithm analyzes the PF charged hadron candidates in a given jet, and combines the information of displaced tracks with information on secondary vertices (vertices with a position not compatible with the beam spot, reconstructed using a dedicated algorithm refitting all tracks in the event) associated with the jet using multivariate analysis techniques to identify b-jets. Jets passing the multivariate discriminant threshold are thus *b-tagged*, and the b-tagging working point used in this analysis achieves a true b-jet tagging efficiency of 68% with a light flavor jet mis-tag efficiency of 1% for jet kinematics typical of $t\bar{t}$ events, as measured in data [181]. Events in simulation are reweighted to correct for the residual data-to-simulation discrepancies in the b-

jet tagging efficiency and light jet mis-tag efficiency, which are of order 1-5%, as a function of the jet kinematics.

6.6.5 Tau leptons

Tau leptons decay primarily hadronically (65% of τ decays [9]), with the remainder of decays to electrons or muons with two associated neutrinos. Hadronically decaying tau leptons are reconstructed using the "Hadrons plus strips" algorithm [182]. Reconstruction is seeded by jets, where the PF candidate constituency of each jet is checked for combinations which are consistent with one of four classes of τ decay modes: one charged hadron $h^{\pm}$; one charged hadron plus a reconstructed π^0 decay (a *strip* of PF e/γ objects, as the photons of $\pi^0 \to \gamma\gamma$ have a high probability to convert in the tracker material), $h^{\pm}\pi^0$; one charged hadron plus two strips $h^{\pm}\pi^0\pi^0$; and three charged hadrons $h^{\pm}h^{\pm}h^{\mp}$. If the jet is identified to be compatible with one of these decay modes, and if the p_{T} sum of PF charged hadrons, neutral hadrons, and photons within a radius $\Delta R < 0.5$ of the jet is less than 5 GeV, the jet is tagged as a hadronic tau object in this analysis. The visible energy of the hadronic tau decay is measured within 1-2% [183].

6.6.6 Missing transverse momentum

The missing transverse momentum vector $\vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}$, and its associated magnitude $p_{\mathrm{T}}^{\mathrm{miss}}$, is constructed from the negative vectorial sum of all PF candidates reconstructed in a event,

$$\vec{p}_{\mathrm{T}}^{\,\mathrm{miss}} = -\sum_{i\in\{h^{\pm},h^0,e,\mu,\gamma,h^{\pm,PU}\}} \vec{p}_{\mathrm{T},i}. \tag{6.5}$$

Charged hadrons associated with pileup primary vertices ($h^{\pm,PU}$) are included in the sum due to the fact that typical pileup collisions will have little to no genuine $p_{\mathrm{T}}^{\mathrm{miss}}$,

and hence, excluding the charged energy deposits while retaining the neutral energy deposits from pileup biases the energy balance and creates an anomalous $p_{\mathrm{T}}^{\mathrm{miss}}$ that is found to be more significant than the additional stochastic contribution to $p_{\mathrm{T}}^{\mathrm{miss}}$ due to the energy resolution of these charged hadrons. The effect of jet energy corrections is propagated to $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ for all reconstructed jets with $p_{\mathrm{T}} > 15\,\mathrm{GeV}$ and passing loose identification criteria, by adding the uncalibrated jet $\vec{p}_{\mathrm{T}}$ and then subtracting the calibrated jet $\vec{p}_{\mathrm{T}}$,

$$\text{Corrected } \vec{p}_{\mathrm{T}}^{\mathrm{miss}} = \vec{p}_{\mathrm{T}}^{\mathrm{miss}} + \sum_{j \in \mathrm{jets}} \vec{p}_{\mathrm{T},j} - \vec{p}_{\mathrm{T},j}^{\,corr}. \tag{6.6}$$

The $p_{\mathrm{T}}^{\mathrm{miss}}$ resolution is improved by the propagation of these jet energy corrections [184, 185].

The resolution of $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ is measured in data by analyzing distributions of the recoil variable $\vec{u}_{\mathrm{T}}$ in Z boson events, where recoil is defined by the relation

$$\vec{u}_{\mathrm{T}} + \vec{p}_{\mathrm{T}}^{\mathrm{miss}} + \vec{p}_{\mathrm{T}}(\mathrm{Z}) = 0.$$

The recoil is divided into two components: $u_{||}$ is the magnitude of the recoil vector projected along the axis defined by the Z boson transverse momentum, and $u_{\perp}$ the magnitude of the corresponding orthogonal component in the transverse plane. In events where a Z boson and no genuine $p_{\mathrm{T}}^{\mathrm{miss}}$ is produced, $u_{||}/p_{\mathrm{T}}(\mathrm{Z})$ is expected to be unity, and $u_{\perp}$ zero. The width of the core of the distributions of these variables characterizes the resolution of the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$, as, although the tails of these distributions are populated preferentially by events with true $p_{\mathrm{T}}^{\mathrm{miss}}$ (they are in fact very close to the signal region definitions of this analysis), the core is dominated by the expected events due to the very large cross section for Z boson production relative to that of processes with genuine $p_{\mathrm{T}}^{\mathrm{miss}}$. For the dataset used in this analysis: the response $u_{||}/p_{\mathrm{T}}(\mathrm{Z})$ is within 2% of unity for $p_{\mathrm{T}}(\mathrm{Z}) > 50\,\mathrm{GeV}$, and data and simulation agree

to within a similar magnitude; the resolution of $u_{||}$ is about 25 GeV, with mild $p_{\mathrm{T}}(\mathrm{Z})$ dependence; and the resolution of $u_{\perp}$ is 22 GeV [185]. For all resolution variables, a strong dependence on the number of pileup vertices in the event is observed, where the values reported here are averaged over the pileup distribution of data used in this analysis.

Several sources of anomalous high-$p_{\mathrm{T}}^{\mathrm{miss}}$ events have been identified in data, originating from deficiencies in the detector and reconstruction algorithms, and dedicated $p_{\mathrm{T}}^{\mathrm{miss}}$ filters remove events with these anomalies from consideration. The filters applied in this analysis exclude events where: a significant energy deposit is located in a region of ECAL with faulty readout and the local reconstruction recovery of energy using the trigger primitive readout uses a trigger primitive that is saturated, thus the ECAL energy in this region is undermeasured by an unknown amount; a significant anomalous energy deposit is found in an isolated HCAL channel due to noise from high-voltage discharges in the hybrid photodiodes; a significant anomalous energy deposit is found in a region of HCAL consistent with originating from a single readout box, due to electronics noise; a pattern of energy deposits in the CSCs as well as the calorimeters that is consistent with a high-energy muon traversing the detector in the plane of the beam line, as caused by *beam halo* muons that are produced from collisions of beam particles with material far away from the interaction point; identified edge cases of the particle flow reconstruction where a poor quality and high p_{T} false muon track is promoted to a PF muon or charged hadron; or ECAL superclusters are found to be formed from energy deposits consistent with spurious noise signals in single crystals. These filters are applied also to simulation samples, where, as none of the above anomalies are simulated, the efficiency of filtering is found to be over 99.9% [185].

Chapter 7

Analysis strategy

7.1 Overview

The $\ell\ell + p_{\mathrm{T}}^{\mathrm{miss}}$ final state is formed by selecting events which contain two opposite-charged same-flavor reconstructed light lepton objects (i.e. e^+e^- or $\mu^+\mu^-$), and forming a composite candidate object with the summed momentum vector of these two leptons, which represents the Z boson candidate. Selections are then applied on variables constructed from the kinematics of the Z boson candidate, its constituent leptons, and the $p_{\mathrm{T}}^{\mathrm{miss}}$ of the event, as well as on other primary objects in the event, defining the signal region (SR). Backgrounds in the signal region are estimated via both simulation and control regions (CRs) in data, and systematic uncertainties are assigned to the estimates of both signal and background yields in the SR and the various CRs. A statistical model is formed from these background estimates and their uncertainties, which is then used to quantify the level of agreement between the observed data and the expectation in the background-only and the signal plus background hypotheses.

7.2 Background composition

Several Standard Model (SM) processes can produce the $\ell\ell + p_{\mathrm{T}}^{\mathrm{miss}}$ final state. A brief overview of the relative contribution of these backgrounds in the signal region is given below.

As discussed in Section 2.3, the main irreducible background in this analysis is the SM $\mathrm{ZZ} \to 2\ell 2\nu$ process, which contributes approximately 60% of the total expected background yield after the full selection. As discussed in Section 5.4, the simulation of this process is corrected to the best available computations. The final estimate of this background utilizes a data control region, explained in Section 7.5.2.

The subleading background in this analysis is the $\mathrm{WZ} \to \ell\nu\ell\ell$ process, where the lepton from the W boson decay is not identified either because it fails the lepton identification requirements, or because it falls outside the detector acceptance or kinematic selections. This process contributes approximately 25% of the total expected background yield after the full selection. The kinematic distributions (e.g. p_{T} of the Z boson candidate, $p_{\mathrm{T}}^{\mathrm{miss}}$) are similar to those of the ZZ background, and the final estimate of this background utilizes a data control region, explained in Section 7.5.1.

A significant background in this analysis comes from *nonresonant* processes, where the reconstructed Z boson candidate is not a true Z boson, but rather a chance combination of opposite-charged same-flavor leptons in the same event which happens to have an invariant mass near that of the Z boson. The source of light leptons is almost exclusively from processes involving production of two W bosons, as light leptons originating from misreconstruction or misidentification, e.g. leptons produced from meson decays in jets which manage to pass the isolation requirements, are well suppressed by the identification requirements described in Section 6.6. Two such nonresonant processes each contribute about 5% to the signal region after full

selection: top production, either $t\bar{t}$ pair production or tW associated production, where the top quark decays almost exclusively [9] to a W boson and a b quark; and WW diboson production, where both W bosons decay to light leptons. Small contributions from single top quark events produced via s- and t-channel processes, and $Z \to \tau\tau$ events in which τ leptons decay into light leptons and neutrinos are also considered as part of the nonresonant background (NRB). The $gg \to H(\to WW)$ process has similar topology to continuum WW production but makes a negligible contribution after the full selection. The final estimate of the contribution of this background uses a data control region, explained in Section 7.6.

Drell–Yan (DY) production, $Z \to \ell\ell$, has a very large total cross section, and although it has little to no genuine $p_{\mathrm{T}}^{\mathrm{miss}}$ (only via neutrino production in heavy meson decays), it can produce events with large $p_{\mathrm{T}}^{\mathrm{miss}}$ caused mainly by jet energy mismeasurement and detector acceptance effects. The DY process amounts to approximately 5% of the total background. The final estimate of the contribution of this background uses a data control region, explained in Section 7.7.

Additional background sources include triboson processes, top plus Z boson production, and diboson production where the reconstructed final state differs from the true final state. All of these processes have a small fiducial cross section, and contribute less than 2% of the total background after the full selection. The final estimate of the contribution of these backgrounds is taken directly from simulation, generated as described in Section 5.3. The leading contributions to this background category are: $Z\gamma \to 2\ell\gamma$, $ZZ \to 4\ell$, WZZ (inclusive of final states), WWZ inclusive, ZZZ inclusive, and tZq inclusive.

The SM process $Z(\to \ell\ell)H(\to ZZ \to 4\nu)$, which is a subset of the ZH(inv.) signal and accounts for 0.1% of SM Higgs boson decays, is not considered as a background due to its negligible contribution to the signal region.

7.3 Event collection

Events are collected using a combination of several trigger paths to maximize the trigger efficiency. For $Z \to ee$ events, a dielectron trigger with $p_T > 23\,GeV$ for the leading electron and $p_T > 12\,GeV$ for the subleading electron captures the majority of events. A combination of single-electron triggers recovers residual inefficiencies, with p_T thresholds of 25 GeV for electrons in $|\eta| < 2.1$ and 27 GeV for electrons in the full η acceptance. All of the above trigger paths include isolation requirements, so to recover any inefficiency due to deficiencies in the L1 or HLT isolation algorithms, and also for the case where the Z boson has a very high momentum and hence its daughter electrons are collimated and spoil each others' isolations, a single electron trigger with no isolation requirements and a p_T threshold of 115 GeV is also employed. For $Z \to \mu\mu$ events, an isolated dimuon trigger with thresholds of $p_T > 17\,GeV$ for the leading and 8 GeV for the subleading muon captures the majority of events. An isolated single muon trigger with $p_T > 24\,GeV$, as well as one with no isolation requirement and a p_T threshold of 50 GeV, is used to recover residual inefficiencies in the collection of events.

The efficiency of this trigger strategy is measured in data by utilizing a *reference trigger*. Events are collected using a prescaled single lepton *reference* trigger with relaxed requirements, such that it has a very high efficiency with respect to events selected by the trigger combination. In this way, the efficiency of the trigger combination can be evaluated with respect to that of the reference trigger with minimal bias. The reference trigger efficiency can be measured separately using the tag and probe technique (Section 6.6.1). The efficiency of the trigger combination can then be expressed

$$\epsilon_{\text{trig.}}(p_T^{\ell 1}, p_T^{\ell 2}) = \left[1 - (1 - \epsilon_{\text{ref}}(p_T^{\ell 1}))(1 - \epsilon_{\text{ref}}(p_T^{\ell 2}))\right] \cdot \epsilon_{\text{comb.}|\text{ref}}(p_T^{\ell 1}, p_T^{\ell 2}), \tag{7.1}$$

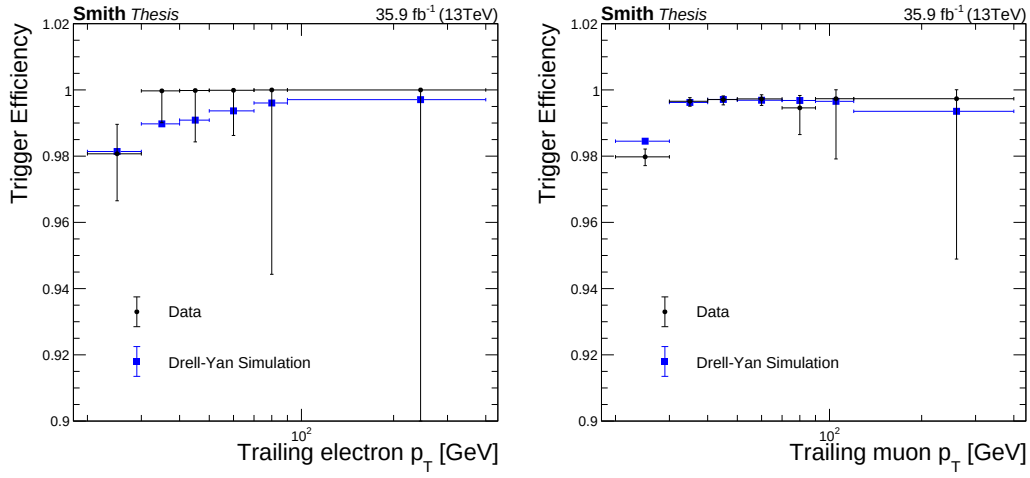

Figure 7.1: Trigger efficiency in data and simulation, as a function of subleading lepton p_{T}, for dielectron (left) and dimuon (right) candidates selected in this analysis, with $|m_{\ell\ell} - m_{\mathrm{Z}}| < 15\,\mathrm{GeV}$ and $p_{\mathrm{T}}^{\ell\ell} > 60\,\mathrm{GeV}$.

where $p_{\mathrm{T}}^{\ell 1}$ and $p_{\mathrm{T}}^{\ell 2}$ are the transverse momenta of the reconstructed leading and subleading lepton, respectively. The efficiency calculation can be made differential also in lepton pseudorapidity, as inefficiencies are typically localized near detector element boundaries; however, as the reference triggers are typically heavily prescaled, there is not sufficient data statistics to effectively probe these effects. The trigger efficiency in simulation is evaluated directly via the second term of Eqn. 7.1, as no observational bias in the event collection is present in simulation. The measured efficiency in data and simulation is shown for dielectron and dimuon events in Fig. 7.1 as a function of $p_{\mathrm{T}}^{\ell 2}$ (averaged over $p_{\mathrm{T}}^{\ell 1}$), with respect to the selection of a reconstructed dilepton pair with invariant mass within 15 GeV of the Z boson mass ($m_{\mathrm{Z}} = 91.1876\,\mathrm{GeV}$ [9]) and a transverse momentum $p_{\mathrm{T}}^{\ell\ell} > 60\,\mathrm{GeV}$, which is a subset of the final selection requirements of this analysis. The simulation is found to model the trigger efficiency correctly, within the uncertainty of this measurement.

7.4 Event selection

Events are required to have exactly two ($N_\ell = 2$) well-identified, isolated leptons of the same flavor and opposite electric charge (e^+e^- or $\mu^+\mu^-$). The leading electron (muon) of the pair must have $p_\mathrm{T} > 25$ (20) GeV, while $p_\mathrm{T} > 20$ GeV is required for the subleading lepton, so that the reconstructed objects have a transverse momentum a bit higher than the trigger threshold to avoid effects of trigger object p_T resolution. Distributions of $m_{\ell\ell}$ and $p_\mathrm{T}^{\ell\ell}$ are shown in Fig. 7.2 with only the lepton kinematic selections applied. In Fig. 7.3, the same distributions are shown after the final selections are applied in the analysis but ignoring the selection requirement of the variable shown, called $n-1$ *distributions*, which show the exclusive effect of these selections. The $n-1$ distributions shown in this section are prior to the maximum likelihood fit (Section 7.8), and uncertainties in these distributions represent only the finite statistics of both data and simulation, ignoring systematic effects (Section 7.9). Hence, only qualitative statements about the level of agreement between data and simulation in these distributions can be made. The dilepton invariant mass ($m_{\ell\ell}$) is required to be within 15 GeV of the established Z boson mass m_Z, which rejects the bulk of the nonresonant background. The dilepton transverse momentum ($p_\mathrm{T}^{\ell\ell}$) must be larger than 60 GeV, which rejects the bulk of the $\mathrm{Z} \to \ell\ell$ background initially, but has no exclusive effect after the combination of the $p_\mathrm{T}^\mathrm{miss}$ and $p_\mathrm{T}^\mathrm{miss}$ balance selections, which are discussed later in this section.

Little hadronic activity is expected in the signal events because to leading order (LO) the final state does not contain any jets. To remove backgrounds with jets in the final state, such as $\mathrm{t\bar{t}}$ production, events having more than one jet with $p_\mathrm{T} > 30$ GeV are rejected. The top quark background is further suppressed by applying a b jet veto: events with at least one b-tagged jet (as described in Section 6.6.4) with

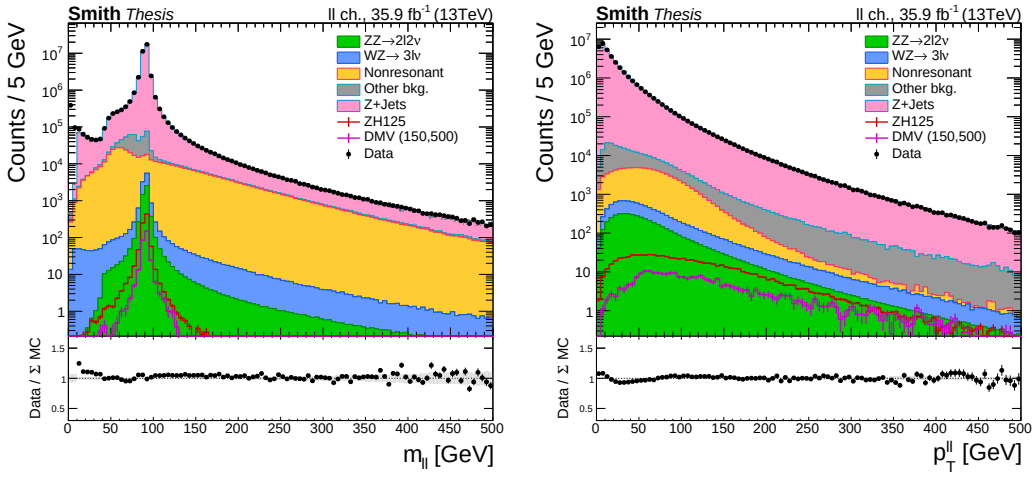

Figure 7.2: Reconstructed Z boson candidate invariant mass ($m_{\ell\ell}$, left) and transverse momentum after the selection $|m_{\ell\ell} - m_{\mathrm{Z}}| < 15\,\mathrm{GeV}$ (right) for the sum of dielectron and dimuon events collected in data as well as the expected event counts from simulation. Selections on the lepton candidates are as described in the text. Error bands indicate statistical uncertainties only.

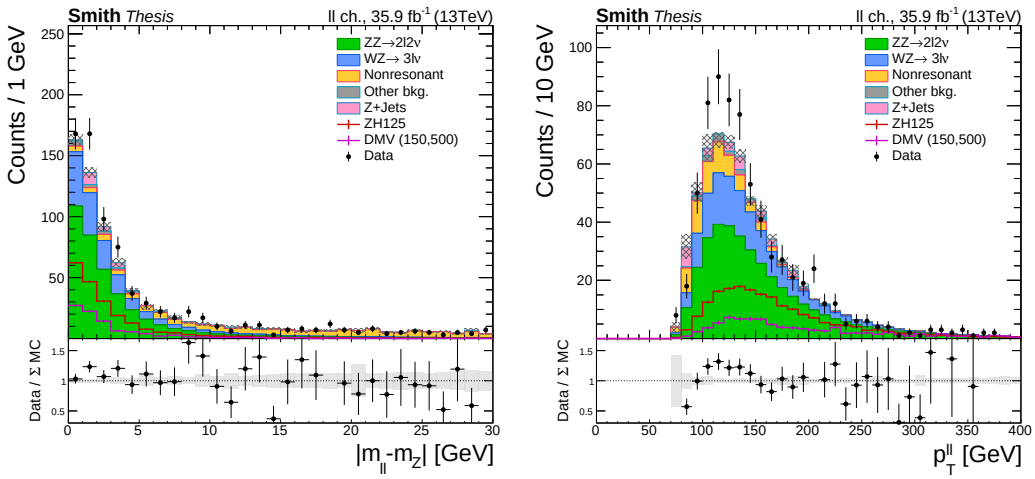

Figure 7.3: Distributions showing the exclusive effect of the selections $|m_{\ell\ell} - m_{\mathrm{Z}}| < 15\,\mathrm{GeV}$ (left) and $p_{\mathrm{T}}^{\ell\ell} > 60\,\mathrm{GeV}$ (right), after all other selections are applied, for the sum of dielectron and dimuon events collected in data as well as the expected event counts from simulation. Error bands indicate statistical uncertainties only.

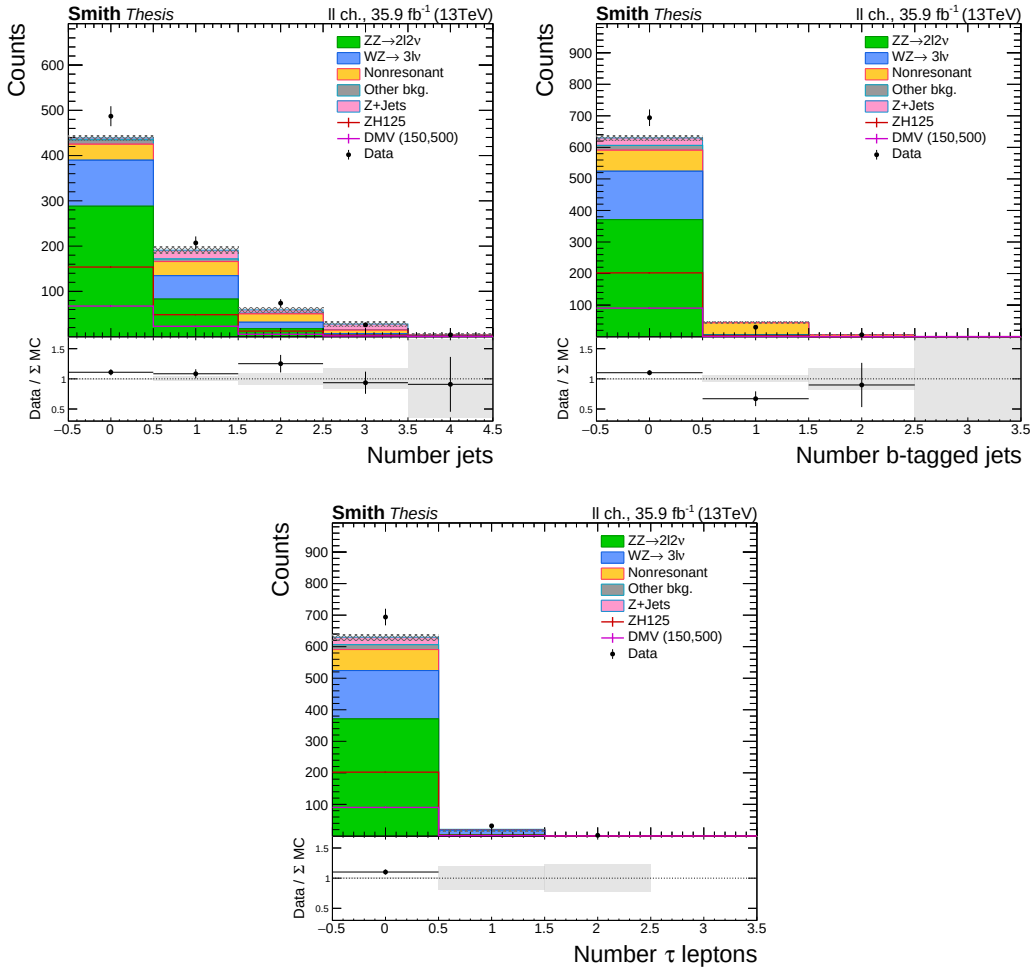

Figure 7.4: Distributions showing the exclusive effect of the selections on the number of jets (less than 2, top left), number of b-tagged jets (less than 1, top right), and number of τ_h candidates (less than 1, bottom) after all other selections are applied, for the sum of dielectron and dimuon events collected in data as well as the expected event counts from simulation. Error bands indicate statistical uncertainties only.

$p_T > 20\,\text{GeV}$ reconstructed within the tracker acceptance, $|\eta| < 2.4$, are removed. To reduce the WZ background in which both bosons decay leptonically, as well as other backgrounds containing extra leptons (e.g. $ZZ \to 4\ell$), events containing additional electrons (muons) with $p_T > 10$ (5) GeV and events with hadronically decaying τ leptons (τ_h) with $p_T > 18\,\text{GeV}$ are removed. Fig. 7.4 shows the jet, b-tagged jet, and τ multiplicity $n-1$ distributions.

The event selection is then optimized using three variables: the $p_{\mathrm{T}}^{\mathrm{miss}}$, the azimuthal angle formed between the dilepton p_{T} and the missing transverse momentum vector, $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell\ell}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}})$, and the $p_{\mathrm{T}}^{\mathrm{miss}}$-$p_{\mathrm{T}}^{\ell\ell}$ balance ratio, $|p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\ell\ell}|/p_{\mathrm{T}}^{\ell\ell}$. These variables are powerful in suppressing reducible background processes, such as DY and top quark production. The selection criteria applied to these variables are optimized in order to obtain the best expected signal sensitivity for a wide range of DM parameters that are considered. For each possible set of selections, the full analysis is repeated, including the estimation of backgrounds from control samples in data and the systematic uncertainties. The final selection criteria obtained after optimization are: $p_{\mathrm{T}}^{\mathrm{miss}} > 100\,\mathrm{GeV}$, $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell\ell}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}) > 2.6\,\mathrm{rad}$, and $|p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\ell\ell}|/p_{\mathrm{T}}^{\ell\ell} < 0.4$. Fig. 7.5 shows the $n-1$ distributions of these variables.

To avoid positive biases in the $p_{\mathrm{T}}^{\mathrm{miss}}$ calculation due to jet mismeasurement, in events with one jet a threshold is applied on the azimuthal angle between this jet and the missing transverse momentum, $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,j}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}) > 0.5\,\mathrm{rad}$. To reduce the contribution from nonresonant backgrounds such as WW and $\mathrm{t\bar{t}}$, a requirement on the distance between the two leptons in the (η, ϕ) plane, $\Delta R_{\ell\ell} < 1.8$, is applied. Fig. 7.6 shows the $n-1$ distributions of these variables.

The signal extraction strategy is based on a maximum likelihood fit (Section 7.8) to the $p_{\mathrm{T}}^{\mathrm{miss}}$ spectrum in data after applying the above selection criteria defining the signal region. The $p_{\mathrm{T}}^{\mathrm{miss}}$ variable shows the strongest signal separation from the background, with $p_{\mathrm{T}}^{\ell\ell}$ the second most powerful variable, while the other selection variables give similar distributions for signal and irreducible background processes, as can be inferred from the $n-1$ distribution figures. A summary of the selection criteria for the SR is given in Table 7.1. The number of data events remaining after each selection, as well as the expected event counts from simulation, are listed in Table 7.2.

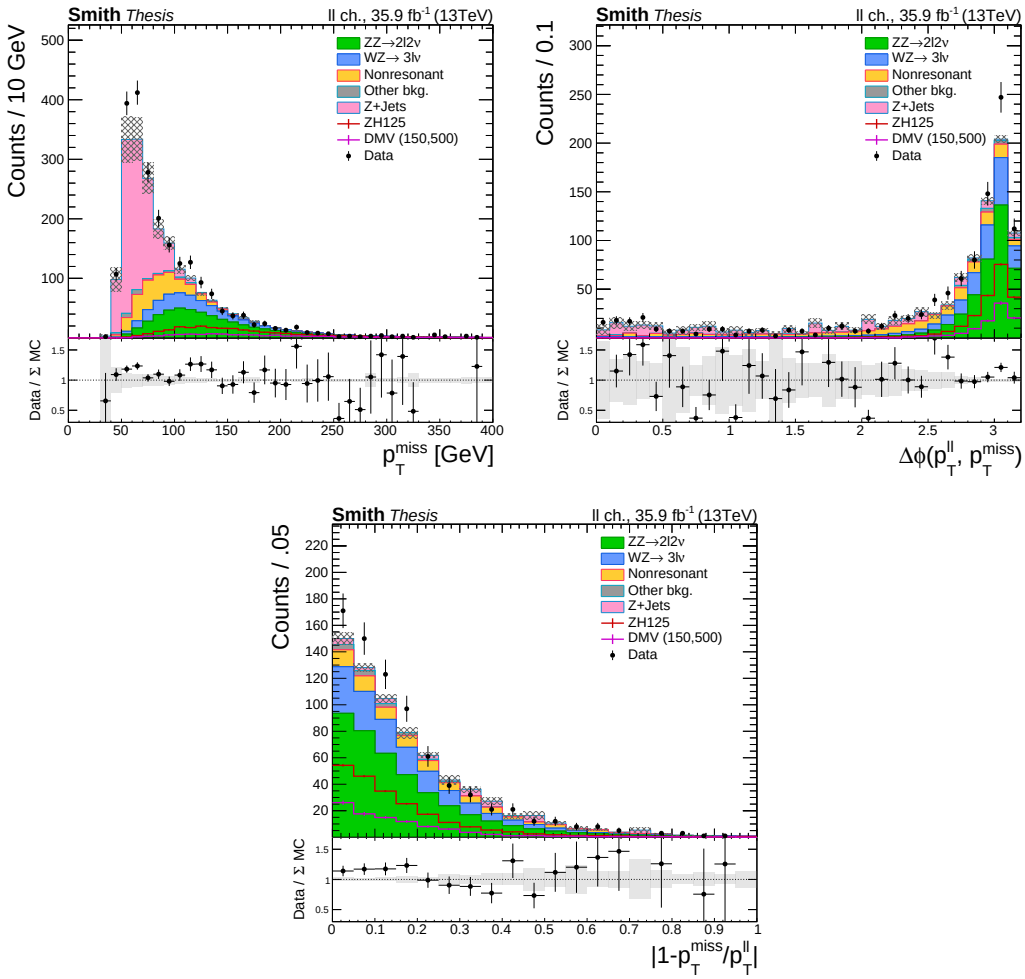

Figure 7.5: Distributions showing the exclusive effect of the selections $p_{\mathrm{T}}^{\mathrm{miss}} > 100\,\mathrm{GeV}$ (top left), $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell\ell}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}) > 2.6\,\mathrm{rad}$ (top right), and $|p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\ell\ell}|/p_{\mathrm{T}}^{\ell\ell} < 0.4$ (bottom) after all other selections are applied, for the sum of dielectron and dimuon events collected in data as well as the expected event counts from simulation. Error bands indicate statistical uncertainties only.

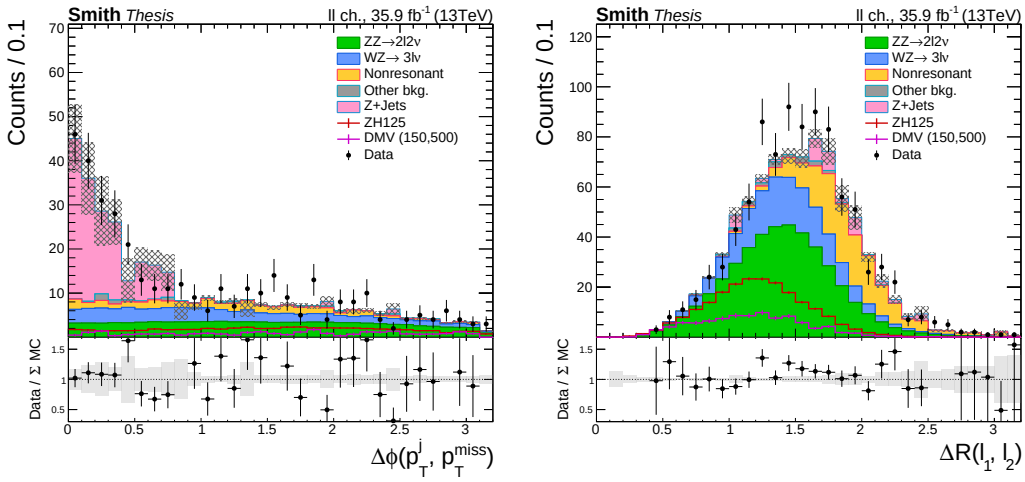

Figure 7.6: Distributions showing the exclusive effect of the selections $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,j}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}) > 0.5$ and $\Delta R_{\ell\ell} < 1.8$ after all other selections are applied, for the sum of dielectron and dimuon events collected in data as well as the expected event counts from simulation. Error bands indicate statistical uncertainties only.

Selection	Requirement	Reject
N_ℓ	$=2$	WZ, triboson
p_{T}^{ℓ}	$>25/20\,\mathrm{GeV}$ for electrons $>20\,\mathrm{GeV}$ for muons	QCD
Z boson mass	$\lvert m_{\ell\ell} - m_{\mathrm{Z}}\rvert < 15\,\mathrm{GeV}$	WW, top quark
$p_{\mathrm{T}}^{\ell\ell}$	$>60\,\mathrm{GeV}$	DY
Jet counting	$\leq$1 jet with $p_{\mathrm{T}}^{j} > 30\,\mathrm{GeV}$	DY, top quark, triboson
b tagging veto	0 b-tagged jets with $p_{\mathrm{T}}^{j} > 20\,\mathrm{GeV}$	Top quark, triboson
τ lepton veto	0 τ_{h} cand. with $p_{\mathrm{T}}^{\tau} > 18\,\mathrm{GeV}$	WZ
$p_{\mathrm{T}}^{\mathrm{miss}}$	$>100\,\mathrm{GeV}$	DY, WW, top quark
$\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell\ell}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}})$	$>2.6\,\mathrm{rad}$	DY
$\lvert p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\ell\ell}\rvert / p_{\mathrm{T}}^{\ell\ell}$	<0.4	DY
$\Delta\phi(\vec{p}_{\mathrm{T}}^{\,j}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}})$	$>0.5\,\mathrm{rad}$	DY, WZ
$\Delta R_{\ell\ell}$	<1.8	WW, top quark

Table 7.1: Summary of the kinematic selections for the signal region.

7.5 Diboson background estimation

The ZZ and WZ processes contribute to the SR via the ZZ $\to \ell\ell\nu\nu$ and WZ $\to \ell\nu\ell\ell$ decay modes, respectively, where the decay products of one boson are not detected. The background estimate for these processes is improved by selecting control regions with alternative decay modes that not only provide a normalization based on CRs in data, but also probe the p_{T} distribution of this lost boson via measurement of the p_{T} of a proxy boson in the CR where its decay products are detected, as the p_{T} distribution of the boson is expected to be independent of the decay mode. In this way, the $p_{\mathrm{T}}^{\mathrm{miss}}$ spectra of these processes are constrained with respect to their theoretical predictions.

The ability of the simulation to correctly model the lost-boson rapidity is important, as the SR rapidity acceptance of the lost boson is necessarily larger than the rapidity acceptance of the proxy boson in each CR, due to the fact that the visible decay products of the proxy boson in the CR must be inside the detector acceptance. The impact of possible data-to-simulation discrepancies in the high-rapidity portion of diboson background in the SR is suppressed by the fact that, as measured in simulation, the majority of the WZ and ZZ contamination in the SR is comprised of events where the lost boson is within the rapidity range of the CRs. The transverse momentum and rapidity distributions of the lost boson in the SR and of the proxy boson in the corresponding CR are shown for simulated WZ and ZZ events in Fig. 7.7. In addition, the proxy boson rapidity distributions in the CRs (or its visible lepton, in the case of the WZ CR) show a good agreement between data and simulation, as shown in Figs. 7.8 and 7.10.

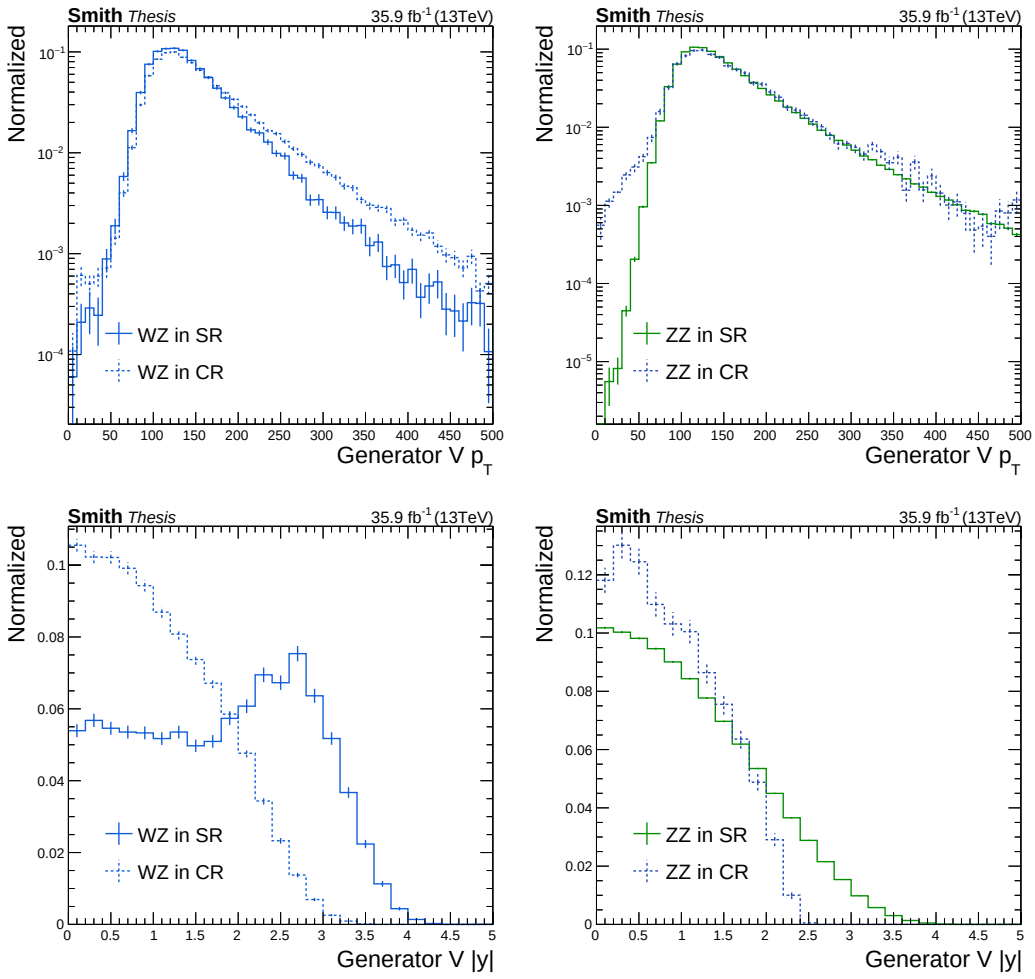

Figure 7.7: Transverse momentum (p_{T}, top) and rapidity ($|y|$, bottom) distributions of the lost boson in the signal region and of the proxy boson in the corresponding control region for simulated WZ events (left) and ZZ events (right), showing that the control regions probe a kinematic phase space similar to that of the corresponding background process in the signal region.

7.5.1 The WZ control region

The WZ control region is formed from events with three well-reconstructed charged leptons. In this case, the CR is populated by events with the same decay mode as the SR, but no leptons are lost to identification or acceptance requirements. A Z boson candidate is selected in the same manner as for the SR, and an additional electron or muon, with identical quality requirements as applied to the leptons in the SR, is required. To enhance the purity of the WZ selection, $p_{\mathrm{T}}^{\mathrm{miss}}$ of at least 30 GeV is required, the invariant mass of three leptons is required to be larger than 100 GeV, and the invariant masses of all opposite-sign, same-flavor lepton pairs are required to be larger than 4 GeV. Backgrounds in this CR are similar to those in the SR, dominated by nonprompt background from the $\mathrm{Z} + jets$ process where a jet is misidentified as a lepton (typically that of the W boson decay due to the Z boson mass constraint), which accounts for about 10% of the CR yield, and the $\mathrm{Z}(\ell\ell)\gamma$ process where the photon is misidentified as an electron (primarily via conversions), which accounts for about 5% of the CR yield. All background estimates for this CR are taken from simulation.

The W boson p_{T} ("emulated $p_{\mathrm{T}}^{\mathrm{miss}}$") is estimated by calculating the vectorial sum of the $\vec{p}_{\mathrm{T}}^{\,\mathrm{miss}}$ vector and the transverse momentum vector ($\vec{p}_{\mathrm{T}}$) of the third charged lepton. In simulation, the majority (over 70%) of WZ background contamination in the signal region originates from events where over 90% of the W boson transverse momentum is carried by one or more neutrinos from the W boson decay. Thus, the majority of the W boson rapidity distribution in the SR is central, although it is less central than in the WZ CR. Neither the SR nor the WZ CR topology can probe the W boson rapidity directly. However, for the WZ CR, good agreement between data and simulation in the third lepton pseudorapidity distributions is observed, as shown in Fig. 7.8, left.

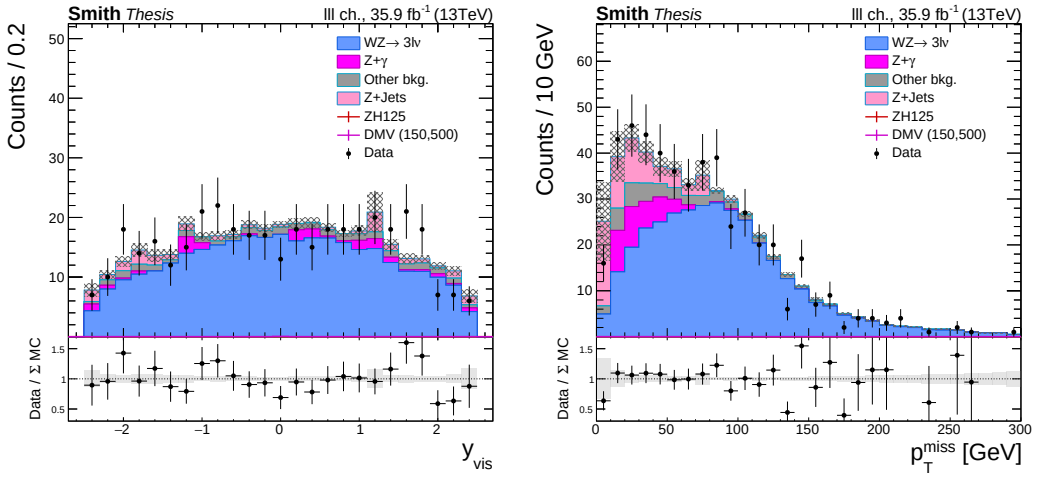

Figure 7.8: Pseudorapidity of the third lepton (left) and $p_{\mathrm{T}}^{\mathrm{miss}}$ (right) distributions in the WZ control region, for events collected in data as well as the expected event counts from simulation. For the $p_{\mathrm{T}}^{\mathrm{miss}}$ distribution, all control region selections are applied other than the requirement that $p_{\mathrm{T}}^{\mathrm{miss}} > 30\,\mathrm{GeV}$. Error bands indicate statistical uncertainties only.

A minor source of WZ background contamination in the SR originates from events where the visible lepton from a W boson decay failed identification requirements. Data-to-simulation discrepancies in this contribution would also manifest in the measured WZ CR $p_{\mathrm{T}}^{\mathrm{miss}}$ distributions, for which no such mismodeling effects are evident, as shown in Fig. 7.8, right.

Using the emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ in place of the reconstructed $p_{\mathrm{T}}^{\mathrm{miss}}$, the same selection is applied as for the SR. However, since there is no danger of CR contamination from WZ $\to \tau\nu\ell\ell$ or top quark backgrounds, no veto on additional τ_{h} or b jet candidates is applied. The resulting emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ spectrum is shown in Fig. 7.9 (top left), showing good agreement between data and simulation. No data is observed in this CR for emulated $p_{\mathrm{T}}^{\mathrm{miss}} > 500\,\mathrm{GeV}$, which is compatible ($p = 0.17$, considering only statistical uncertainties) with the expectation from simulation of 1.8 events.

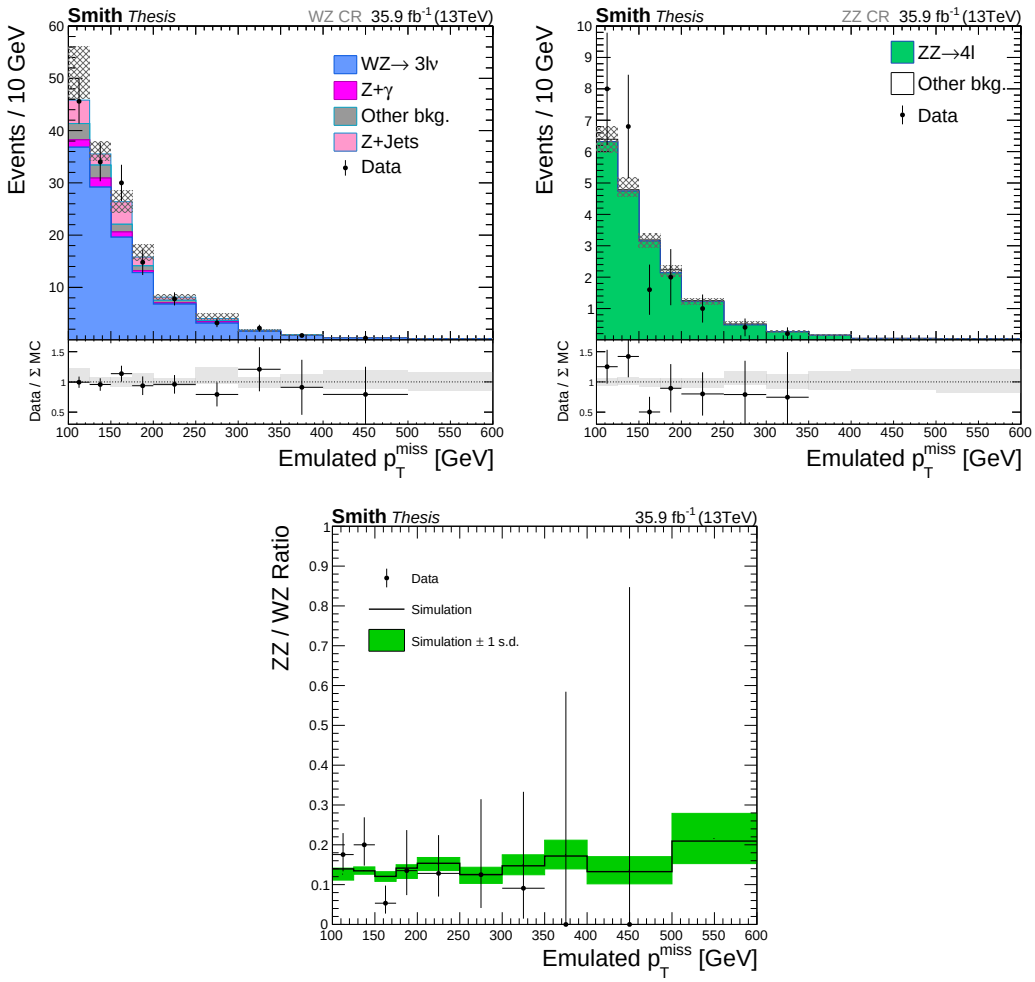

Figure 7.9: Emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ distributions in data and simulation for the WZ $\rightarrow 3\ell\nu$ (top left) and ZZ $\rightarrow 4\ell$ (top right) control regions, and the ratio between both distributions (bottom). No events are observed with emulated $p_{\mathrm{T}}^{\mathrm{miss}} > 500\,\mathrm{GeV}$ in either control region. Uncertainty bands correspond to the combined statistical and systematic components.

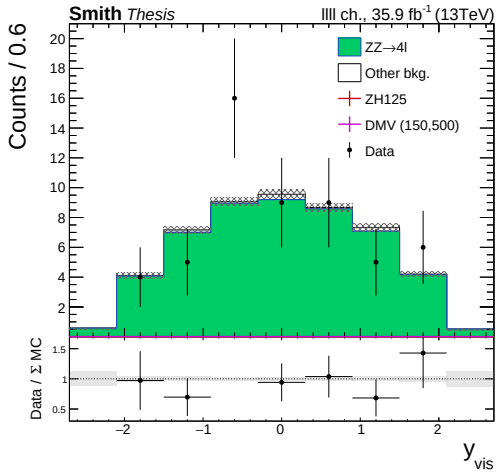

Figure 7.10: Rapidity distributions of the proxy Z boson in the ZZ control region, for events collected in data as well as the expected event counts from simulation. Error bands indicate statistical uncertainties only.

7.5.2 The ZZ control region

The ZZ control region is formed from events with four well-reconstructed charged leptons. In addition to a signal-like Z boson candidate, a second Z boson candidate is required, the constituents of which only need to pass relaxed lepton quality requirements, namely those used to veto additional leptons in the SR. This choice reflects the very high purity of the four-lepton selection. For both Z boson candidates, the same mass constraint as in the SR is applied. Backgrounds, dominated by triboson processes, are almost negligible in this CR and are taken from simulation.

Similar to the WZ case, the emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ is calculated as the vectorial sum of the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ and the $\vec{p}_{\mathrm{T}}$ of the Z boson with the larger mass difference to the nominal value of m_{Z} of the two Z bosons identified in the event. The choice of which Z boson to use as a proxy for an invisibly decaying one does not significantly alter the emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ spectrum. In this CR, the rapidity of the proxy boson is observable, for which good agreement between data and simulation is found, as shown in Fig. 7.10.

The same selection as in the SR is then applied using the emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ in place of the reconstructed $p_{\mathrm{T}}^{\mathrm{miss}}$, with the exception of the τ lepton and b jet vetoes. The resulting emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ spectrum is shown in Fig. 7.9 (top right), showing good agreement between data and simulation. No data is observed in this CR for emulated $p_{\mathrm{T}}^{\mathrm{miss}} > 350\,\mathrm{GeV}$, which is compatible ($p = 0.2$, considering only statistical uncertainties) with the expectation from simulation of 1.6 events.

7.5.3 VV differential distributions

Due to a limited event count in the ZZ control region, the normalizations of the WZ and ZZ processes in the WZ and ZZ CRs and the SR are controlled by a single free parameter (μ_{VV} in Eqn. 7.2) in the maximum likelihood fit, with their relative normalizations fixed by the theoretical predictions for the WZ and ZZ processes in each $p_{\mathrm{T}}^{\mathrm{miss}}$ bin. The predictions for these processes are obtained from fully reconstructed simulated events generated as described in Chapter 5 with the higher-order corrections described in Section 5.4 applied.

The product of the magnitudes of the NLO EW corrections and the inclusive NLO QCD corrections [136] are used as an estimate of the missing NLO EW $\times$ NLO QCD contribution, which is not used as a correction, but rather assigned as an uncertainty. The uncertainties in the WZ and ZZ EW corrections are assumed to be anticorrelated as a conservative measure. The uncertainty associated with the NNLO QCD corrections for both processes is represented by the QCD scale variation uncertainties evaluated on the NLO QCD simulation sample for the respective process, as described in Section 7.9. Figure 7.9 (bottom) shows the ratio of ZZ to WZ CR yields per emulated $p_{\mathrm{T}}^{\mathrm{miss}}$ bin, which probes the validity of taking the relative normalizations from simulation. Good agreement is observed between data and simulation.

7.6 Nonresonant background estimation

The contribution of the nonresonant backgrounds is estimated from a CR based on events with two leptons of different flavor ($e^{\pm}\mu^{\mp}$) that pass all other analysis selections. Nonresonant background consists mainly of leptonic W boson decays in $\mathrm{t\bar{t}}$, tW, and WW events, where the dilepton mass happens to fall inside the Z boson mass window. Small contributions from single top quark events produced via s- and t-channel processes, and $\mathrm{Z} \to \tau\tau$ events in which τ leptons decay into light leptons and neutrinos are also considered in the NRB estimation.

The method assumes lepton flavor symmetry in the final states of these processes. Since the leptonic decay branching fraction to the ee, $\mu\mu$, and $e\mu$ final states from NRB are 1:1:2, the $e\mu$ events selected inside the Z boson mass window can be extrapolated to the ee and $\mu\mu$ channels. To account for differences in efficiency for electrons and muons, a correction factor k_{ee} is derived by comparing the NRB yields for the ee and $\mu\mu$ channels:

$$k_{ee} = \frac{\epsilon_e}{\epsilon_\mu} = \sqrt{\frac{N^{ee}_{\mathrm{NRB}}}{N^{\mu\mu}_{\mathrm{NRB}}}}$$

under the assumption that there are no efficiency correlations between the two leptons. In simulation, k_{ee} is found to be about 0.88 for the final selection. With this correction factor, the relation between the NRB yields in the SR and CR is:

$$N^{\ell\ell}_{\mathrm{NRB}} = \frac{1}{2}\left(k_{ee} + \frac{1}{k_{ee}}\right) N^{e\mu}_{\mathrm{NRB}}.$$

The ratio of the NRB contributions in the SR and CR is fixed by this relation. Their normalization is controlled by a common scaling parameter that is left to float in the maximum likelihood fit. Perturbations in the predicted transfer factor due to data-to-simulation discrepancies in k_{ee} are suppressed upon summing the $ee + \mu\mu$ channels. The uncertainty in the transfer factor is set conservatively to 20%, based on the

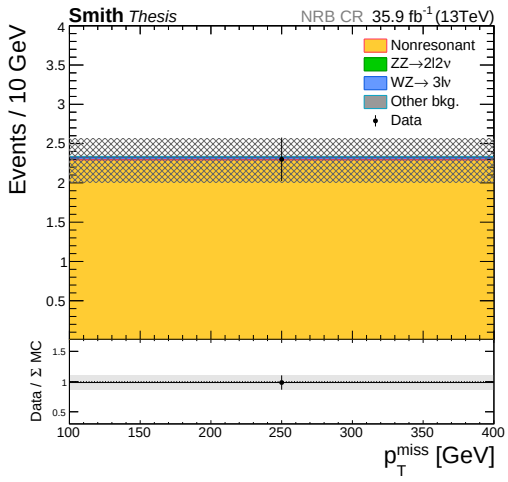

Figure 7.11: Observed and expected yields of $e\mu$ events in the nonresonant background control region. Uncertainty bands correspond to the combined statistical and systematic components.

observed variation in its value in simulation as each successive selection defining the SR is applied. The observed and expected yields of $e\mu$ events is shown in Fig. 7.11.

7.7 Drell–Yan background estimation

The DY background is dominant in the region of low $p_{\mathrm{T}}^{\mathrm{miss}}$. This process does not typically produce undetectable particles with significant transverse momentum, apart from rare semileptonic decays of heavy mesons in jets that may produce high p_{T} neutrinos, therefore any nonzero $p_{\mathrm{T}}^{\mathrm{miss}}$ arises from the limited detector acceptance and mismeasurement.

The estimation of this background uses simulated DY events, for which the normalization is taken from data in a sideband CR of $50 \leq p_{\mathrm{T}}^{\mathrm{miss}} \leq 100\,\mathrm{GeV}$, with all other selections applied. Two control regions (where a larger DY background contribution is expected) are formed by omitting the selection on $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,j}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}})$ and, respectively, inverting the selections on $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell\ell}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}})$ and on the $p_{\mathrm{T}}^{\mathrm{miss}}$ balance vari-

able. For each control region, the distribution of the ratio of the nominal particle flow $p_{\mathrm{T}}^{\mathrm{miss}}$ (Section 6.6.6) and a $p_{\mathrm{T}}^{\mathrm{miss}}$ variable constructed only from calorimeter information[1], $|p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\mathrm{miss,calo}}|/p_{\mathrm{T}}^{\mathrm{miss}}$, is shown in Fig. 7.12 (left). This variable is sensitive to mis-modeling of the charged to neutral hadron ratio in jets and the underlying event, as well as the response of particle flow reconstruction to this ratio. Additionally, for events containing a jet, the distribution of $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,j}, \vec{p}_{\mathrm{T}}^{\,\mathrm{miss}})$ is shown in Fig. 7.12 (right). In these control regions, the simulation is found to model the data well.

The sideband CR is included in the maximum likelihood fit, for which the normalization factor is found to be consistent with unity, and a 100% uncertainty is assigned to the resulting DY estimate in order to cover the extrapolation from this CR to the SR. This uncertainty has little effect on the results owing to the small overall contribution from the DY process in the high-$p_{\mathrm{T}}^{\mathrm{miss}}$ SR of this analysis.

[1] Muon tracks are also considered in calo $p_{\mathrm{T}}^{\mathrm{miss}}$

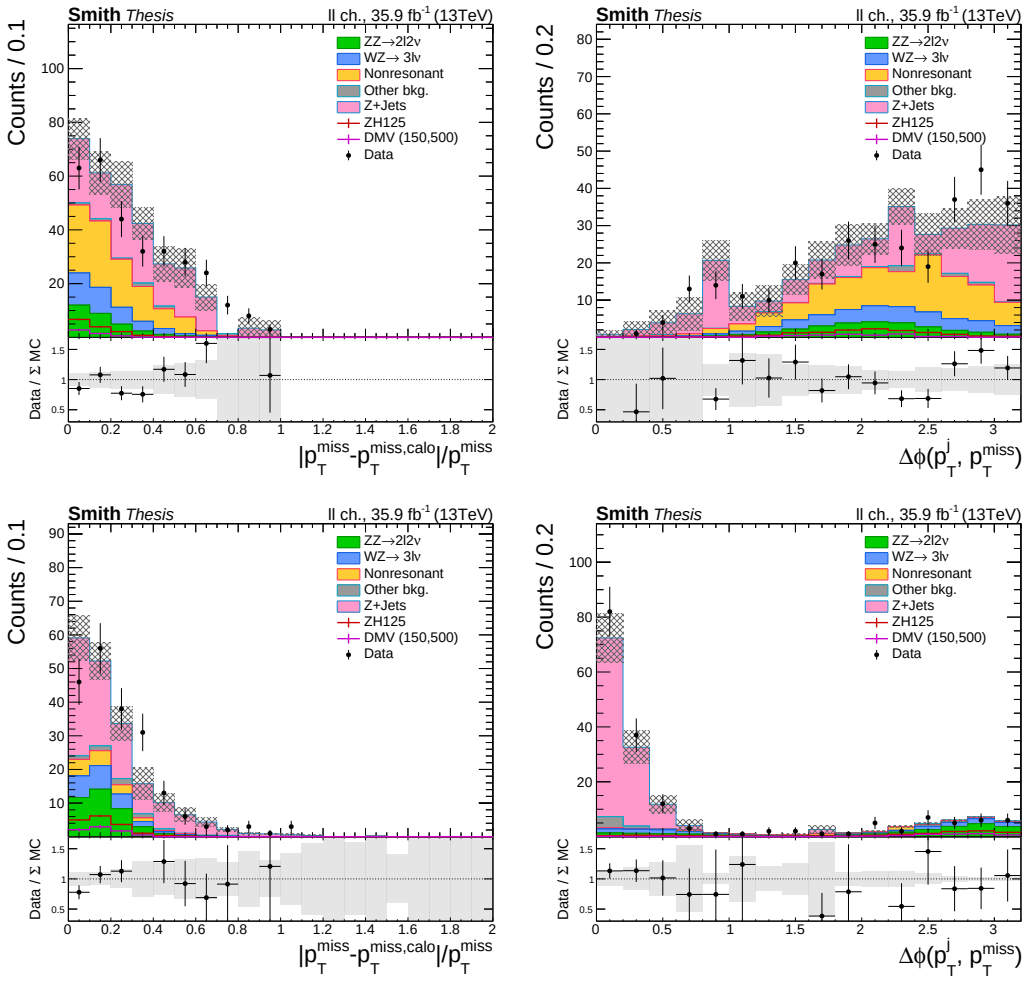

Figure 7.12: Distributions of the ratio of the nominal particle flow $p_{\mathrm{T}}^{\mathrm{miss}}$ and a $p_{\mathrm{T}}^{\mathrm{miss}}$ variable constructed only from calorimeter information ($|p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\mathrm{miss,calo}}|/p_{\mathrm{T}}^{\mathrm{miss}}$, left) and of $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,j}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ (right), for two control regions constructed from the SR by omitting the selection on $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,j}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ and inverting the selection on $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell\ell}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ (top) or on $|p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\ell\ell}|/p_{\mathrm{T}}^{\ell\ell}$ (bottom), for events collected in data as well as the expected event counts from simulation. Error bands indicate statistical uncertainties only.

Selection	Event yield								
	ZZ	WZ	NRB	Other	DY	Total bkg.	ZH(inv.)	Vector DM	Data
e^+e^- or $\mu^+\mu^-$	5269.5 ± 1.6	7663.2 ± 8.4	430700 ± 330	307240 ± 340	37017000 ± 13000	37768000 ± 13000	861.3 ± 3.9	295.3 ± 5.4	38879200 ± 6200
Z boson mass	4868.0 ± 1.5	5904.2 ± 7.4	78210 ± 100	172090 ± 250	33941000 ± 12000	34202000 ± 12000	816.9 ± 3.8	281.8 ± 5.3	34788900 ± 5900
$p_{\mathrm{T}}^{\ell\ell}$	1968.39 ± 0.97	2569.6 ± 4.9	35866 ± 53	48520 ± 130	2321900 ± 3300	2410800 ± 3300	580.9 ± 3.1	229.8 ± 4.8	2430800 ± 1600
Jet counting	1620.02 ± 0.88	1466.5 ± 3.7	8514 ± 32	19590 ± 90	1376500 ± 2500	1407700 ± 2500	453.7 ± 2.8	184.4 ± 4.2	1446100 ± 1200
b tagging veto	1591.33 ± 0.87	1433.1 ± 3.6	4544 ± 27	18255 ± 87	1290500 ± 2400	1316300 ± 2400	446.8 ± 2.8	182.5 ± 4.2	1361400 ± 1200
τ lepton veto	1572.97 ± 0.87	1258.5 ± 3.4	4463 ± 27	17648 ± 86	1261600 ± 2400	1286500 ± 2400	442.0 ± 2.8	180.4 ± 4.1	1328100 ± 1200
$p_{\mathrm{T}}^{\mathrm{miss}}$	624.67 ± 0.54	332.5 ± 1.7	727.1 ± 9.5	44.2 ± 3.3	771 ± 51	2499 ± 52	278.4 ± 2.2	126.7 ± 3.5	2473 ± 50
$\Delta\phi(\vec{p}_{\mathrm{T}}^{\ell\ell}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$	553.42 ± 0.51	252.5 ± 1.5	348.6 ± 6.6	31.6 ± 2.6	318 ± 30	1504 ± 31	252.2 ± 2.1	114.3 ± 3.4	1602 ± 40
$\lvert p_{\mathrm{T}}^{\mathrm{miss}} - p_{\mathrm{T}}^{\ell\ell} \rvert / p_{\mathrm{T}}^{\ell\ell}$	448.58 ± 0.46	196.9 ± 1.3	176.9 ± 4.7	20.3 ± 2.2	173 ± 21	1015 ± 22	223.0 ± 2.0	100.0 ± 3.2	1107 ± 33
$\Delta\phi(\vec{p}_{\mathrm{T}}^{j}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$	431.80 ± 0.45	179.8 ± 1.3	166.2 ± 4.6	16.5 ± 1.7	38 ± 11	833 ± 12	215.1 ± 1.9	96.0 ± 3.1	910 ± 30
$\Delta R_{\ell\ell}$	370.79 ± 0.42	153.5 ± 1.2	66.6 ± 2.8	15.3 ± 1.6	23.8 ± 8.3	629.9 ± 9.0	202.2 ± 1.9	90.4 ± 3.0	694 ± 26

Table 7.2: The number of events remaining after each selection, applied in the order specified, for data and simulated events. Simulated events are classified according to the background categories as described in Section 7.2. Yields for two example signal models are included: the ZH(inv.) model with $m_{\mathrm{H}} = 125\,\mathrm{GeV}$, and the vector dark matter model, generated with $m_{\mathrm{med}} = 500\,\mathrm{GeV}$ and $m_{\mathrm{DM}} = 50\,\mathrm{GeV}$. Uncertainties include only the statistical component.

7.8 Likelihood model for signal extraction

The signal extraction strategy is based on a binned maximum likelihood fit to the $p_{\mathrm{T}}^{\mathrm{miss}}$ spectrum of data in the SR as well as to the $p_{\mathrm{T}}^{\mathrm{miss}}$ spectrum of data in the WZ and ZZ control regions, and the two single-bin control regions for nonresonant and Drell–Yan background normalizations. There is no significant difference between the dielectron and dimuon channels in terms of signal-to-background ratio, and the systematic uncertainties due to the finite size of the simulation samples would be inflated if the channels were treated separately in the likelihood model, hence both are treated together when obtaining the final results. The likelihood $\mathcal{L}$ is constructed as follows:

$$
\begin{aligned}
\mathcal{L} = \prod_i \mathcal{P}\left(N^{2\ell}_{obs,i} \middle| \mu_{DY} N^{2\ell}_{DY,i}(\boldsymbol{\theta}) + \mu_{NRB} N^{2\ell}_{NRB,i}(\theta) + N^{2\ell}_{other,i}(\boldsymbol{\theta}) \right. \\
\left. + \mu_{VV}(N^{2\ell}_{ZZ,i}(\boldsymbol{\theta}) + N^{2\ell}_{WZ,i}(\boldsymbol{\theta})) + \mu N^{2\ell}_{Sig,i}(\boldsymbol{\theta})\right) \\
\times \prod_i \mathcal{P}\left(N^{3\ell}_{obs,i} \middle| N^{3\ell}_{other,i}(\boldsymbol{\theta}) + \mu_{VV} N^{3\ell}_{WZ,i}(\boldsymbol{\theta})\right) \\
\times \prod_i \mathcal{P}\left(N^{4\ell}_{obs,i} \middle| N^{4\ell}_{other,i}(\boldsymbol{\theta}) + \mu_{VV} N^{4\ell}_{ZZ,i}(\boldsymbol{\theta})\right) \\
\times \mathcal{P}\left(N^{e\mu}_{obs} \middle| \mu_{NRB} N^{e\mu}_{NRB}(\boldsymbol{\theta}) + N^{e\mu}_{other}(\boldsymbol{\theta})\right) \\
\times \mathcal{P}\left(N^{DYsb}_{obs} \middle| \mu_{DY} N^{DYsb}_{DY}(\boldsymbol{\theta}) + \mu_{NRB} N^{DYsb}_{NRB}(\boldsymbol{\theta}) + N^{DYsb}_{other}(\boldsymbol{\theta}) \right. \\
\left. + \mu_{VV}(N^{DYsb}_{ZZ}(\boldsymbol{\theta}) + N^{DYsb}_{WZ}(\boldsymbol{\theta})) + \mu N^{DYsb}_{Sig}(\boldsymbol{\theta})\right) \\
\times e^{-|\boldsymbol{\theta}|^2/2},
\end{aligned}
\tag{7.2}
$$

where $\mathcal{P}(N|\lambda)$ is the Poisson probability, μ is the signal strength, μ_{VV} is the diboson process normalization, μ_{DY} is the Drell-Yan normalization, μ_{NRB} is the nonresonant background normalization, $N_x^{2\ell}$ is the MC prediction for the yield of process x in the signal region, $N_x^{3\ell}$ is the MC prediction for the yield of process x in the WZ control region, $N_x^{4\ell}$ is the MC prediction for the yield of process x in the ZZ control region, $N_x^{e\mu}$ is the MC prediction for the yield of process x in the NRB control region, N_x^{DYsb}

is the MC prediction for the yield of process x in the Drell–Yan sideband control region, and $\boldsymbol{\theta}$ are nuisance parameters of the model representing the systematic uncertainties described in Section 7.9. These nuisance parameters modify either the per-bin predictions or the overall normalization of the processes which are affected by a given uncertainty. In the case of modifications to per-bin predictions, the uncertainty is referred to as a *shape uncertainty*. All nuisance parameters are assigned a prior probability density function, such that the likelihood receives a penalty term if the parameters are shifted away from the nominal value, i.e. $\theta_i \neq 0$. In this analysis, the nuisance parameters represent uncertainties on multiplicative corrections to the simulation, so the priors are chosen to be log-normal distributed. The yield for process x is modified from its nominal value $\hat{N}_x$ as a function of the nuisance parameter θ_i as:

$$N_x(\theta_i) = \hat{N}_x(1 + \Delta_i N_x/\hat{N}_x)^{\theta_i},$$

where $\Delta_i N_x$ is the change in the yield N_x caused by a one standard deviation ($1\,\sigma$) shift of the values of parameters in simulation that are influenced by the uncertainty i. Typically, the ΔN corresponding to positive and negative $1\,\sigma$ shifts are not equal, in which case a second-order continuous function is constructed to smoothly interpolate between the nominal and $\pm 1\,\sigma$ shifted values [186].

The likelihood of Eqn. 7.2 is maximized (or rather, $-\ln \mathcal{L}$ is minimized) with respect to the parameters $\{\mu, \mu_{VV}, \mu_{DY}, \mu_{NRB}, \theta_i\}$ using the MINUIT2 software package [187] in the ROOFIT framework [188]. The resulting yields in simulation are referred to as *post-fit* yields. All figures and tables in this section other than Figs. 7.9 and 7.11 display *pre-fit* yields, i.e. before this minimization is performed. All figures and tables in Chapter 8 display post-fit yields.

This minimization procedure also defines a posterior probability density function for all nuisance parameters (all parameters except μ), from which upper limits on

a given signal model's signal strength can be computed using the CL_s technique [189, 190]. This method can also be applied to calculate the significance of a data excess; however, as will be shown in Chapter 8, this will not be necessary in this analysis. The technique is summarized as follows:

- Define a test statistic q for a given hypothesized signal strength μ using the likelihood function,

$$q_\mu = \begin{cases} \dfrac{\mathcal{L}(\mu, \hat{\hat{\boldsymbol{\theta}}})}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} & \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} \tag{7.3}$$

  where $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ are the values of the signal strength and nuisance parameters (now including $\mu_{VV}, \mu_{DY}, \ldots$) that maximize $\mathcal{L}$, and $\hat{\hat{\boldsymbol{\theta}}}$ are the values of the nuisance parameters that maximize $\mathcal{L}$ for the given μ.

- Construct the probability density function of q_μ for a given μ, $f(q_\mu|\mu)$, by replacing the various N_{obs} quantities of Eqn. 7.2 with those of a pseudo-experiment and evaluate Eqn. 7.3, for an ensemble of pseudo-experiments created by Monte Carlo sampling of the posterior probability density function of the nuisance parameters. This procedure can be computationally expensive, so in this analysis, an analytic formula for $f(q_\mu|\mu)$ is constructed using the Hessian matrix of the likelihood function as computed numerically by MINUIT2, which is a very good approximation provided the observed number of events in each bin is not too small [191]. No significant difference was found between the results computed using these two methods for a subset of the signal models in this analysis.

- Calculate the level of agreement between the observed data and the hypothesized signal strength with a p-value constructed from the test statistic distribu-

tion,

$$P(\mu) = \int_{q_{\mu,\mathrm{obs}}}^{\infty} f(q_\mu|\mu)\, dq_\mu,$$

from which the CL_s value is defined

$$\mathrm{CL}_s(\mu) = \frac{P(\mu)}{P(0)}. \tag{7.4}$$

The value of μ for which $\mathrm{CL}_s = 1 - 0.95$ then corresponds to the signal strength threshold excluded at 95% confidence level.

7.9 Systematic uncertainties

The uncertainties in the yields from missing higher-order QCD corrections in signal as well as ZZ and WZ background differential cross sections are evaluated by independently varying up and down the factorization and renormalization scales (Section 5.2) by a factor of two [123, 144]. The effect of these variations on the yields is between 5 and 10%. For the ZZ and WZ backgrounds, additional uncertainties related to known higher-order corrections are applied, as discussed in Section 7.5.3.

For the Higgs boson signal, the parton density function (PDF) and α_s uncertainties comprise the cross section normalization uncertainties computed by the LHC Higgs cross section working group [192] and the effect on the signal acceptance of varying the PDFs and α_s within their uncertainties [193]. For other signal models, as well as the WZ and ZZ backgrounds, the effects of the PDF and α_s uncertainties in the signal acceptance are taken into account by finding the $1\,\sigma$ coverage interval of the distribution of yields evaluated from 100 Monte Carlo replicas of the NNPDF set, following the PDF4LHC prescription [193]. The PDF and α_s uncertainties on these processes are found to be about 1–2%.

The efficiencies for triggering on, reconstructing, and identifying isolated leptons are obtained from simulation, and corrected with scale factors as discussed previously.

The associated uncertainty is about 1–2% per lepton. An additional 3% uncertainty associated with the WZ $\to \ell\nu\ell\ell$ events, where the reconstructed lepton from the W boson decay fails identification, is also included.

In order to reproduce b tagging efficiencies and light-flavor jet mistag rates observed in data, an event-by-event reweighting using data-to-simulation scale factors is applied to simulated events, as discussed in Section 6.6.4. The uncertainty associated with this procedure is obtained by varying the event-by-event weight by $\pm 1\,\sigma$. The impact on the final yields due to the b tagging efficiency and mistag rate uncertainties is around 1% for both signal and background.

The impacts of the jet energy scale and resolution uncertainties are estimated by shifting reconstructed jet energies in simulation by $\pm 1\,\sigma$, and each is found to have an effect of about 2% on the yields of the simulated processes after all selections are applied. The impacts of the electron and muon energy scales are evaluated in the same manner, and have a similar effect. Uncertainties in the $p_{\mathrm{T}}^{\mathrm{miss}}$ measurement due to the energy resolution of unclustered PF candidates (i.e., those not associated with an electron, muon, or jet) amount to about 2%.

The uncertainty in the expected yields due to the finite size of the MC samples is considered, and is around 1% for the signal and main backgrounds. The simulated MC samples are reweighted to reproduce the pileup conditions observed in data (Fig. 5.2). The uncertainty related to this procedure is obtained by varying the central value of the estimated inelastic cross section by 5% [102], and is found to be below 1%. The uncertainty assigned to the integrated luminosity measurement is 2.5% [113].

The effect of the systematic uncertainties on the shape of the distribution of the discriminating variable ($p_{\mathrm{T}}^{\mathrm{miss}}$) is taken into account by varying the value of the quantity associated with the uncertainty, and observing the resulting variations in the individual bins of $p_{\mathrm{T}}^{\mathrm{miss}}$.

All these sources of uncertainty are summarized in Table 7.3. The combined uncertainty in the signal efficiency and acceptance is estimated to be about 5% and is dominated by the theoretical uncertainty due to missing higher-order corrections and PDF uncertainties. The total uncertainty in the background estimations in the signal region is about 15%, dominated by the theoretical uncertainties in the ZZ and WZ process description.

Source of uncertainty	Effect (%)					Impact on the exp. limit (%)
	Signal	ZZ	WZ	NRB	DY	
* VV EW corrections	–	10	−4	–	–	14
* Renorm./fact. scales, VV	–	9	4	–	–	2
* Renorm./fact. scales, ZH	3.5	–	–	–	–	
* Renorm./fact. scales, DM	5	–	–	–	–	
* PDF, WZ background	–	–	1.5	–	–	
* PDF, ZZ background	–	1.5	–	–	–	
* PDF, Higgs boson signal	1.5	–	–	–	–	
* PDF, DM signal	1–2	–	–	–	–	
* MC sample size, NRB	–	–	–	5	–	1
* MC sample size, DY	–	–	–	–	30	
* MC sample size, ZZ	–	0.1	–	–	–	
* MC sample size, WZ	–	–	2	–	–	
* MC sample size, ZH	1	–	–	–	–	
* MC sample size, DM	3	–	–	–	–	
NRB extrapolation to the SR	–	–	–	20	–	<1
DY extrapolation to the SR	–	–	–	–	100	<1
Lepton efficiency (WZ CR)	–	–	3	–	–	<1
Nonprompt bkg. (WZ CR)	–	–	–	–	30	<1
Integrated luminosity	2.5					<1
* Electron efficiency	1.5					1
* Muon efficiency	1					
* Electron energy scale	1–2					
* Muon energy scale	1–2					
* Jet energy scale	1–3 (typically anticorrelated w/ yield)					
* Jet energy resolution	1 (typically anticorr.)					
* Unclustered energy ($p_{\mathrm{T}}^{\mathrm{miss}}$)	1–4 (typically anticorr.), strong in DY					
* Pileup	1 (typically anticorrelated)					
* b tagging eff. & mistag rate	1					

Table 7.3: Summary of the systematic uncertainties present in this analysis. Each uncertainty represents the variation of the relative yields of the processes in the SR. Each uncertainty is fully correlated across processes to which it contributes, including those processes that are also present in CRs. The symbol "–" indicates that the systematic uncertainty does not contribute or is deemed negligible. For minor backgrounds, systematic uncertainties are omitted because of the smallness of their contribution. For shape uncertainties (indicated with a *), the numbers correspond to the overall effect of the shape variation on the yield or acceptance. The impact on the expected upper limit for the signal strength, i.e., the relative decrease in the median expected upper limit for the signal strength upon removing the nuisance term, is evaluated with respect to the SM H(inv.) signal and presented in the last column.

Chapter 8

Results

8.1 Introduction

The analysis strategy described in Chapter 7 is applied to the dataset collected in 2016 (described in Section 4.11), and the resulting numbers of observed and expected events in the signal region are presented. The procedure of Section 7.8 is used to set exclusion limits on the parameters of the models presented in Chapter 2. These limits are then compared to the relevant limits from other experiments, where available. Finally, a method for re-interpreting the result of this analysis in the context of alternative signal hypotheses is presented.

8.2 Observed data

The numbers of observed and expected events in the signal region after the final selection are shown in Table 8.1. The numbers of expected events for the various backgrounds are those that maximize the likelihood function of Eqn. 7.2 under the background-only ($\mu = 0$) hypothesis. The numbers of expected events for the signals

Process	Signal region yield
qqZH(inv.) $m_{\mathrm{H}} = 125\,\mathrm{GeV}$, $\mathcal{B}(\mathrm{H} \to \mathrm{inv.}) = 1$	159 ± 13
ggZH(inv.) $m_{\mathrm{H}} = 125\,\mathrm{GeV}$, $\mathcal{B}(\mathrm{H} \to \mathrm{inv.}) = 1$	43 ± 11
DM, vector mediator $m_{\mathrm{med}} = 500\,\mathrm{GeV}$, $m_{\mathrm{DM}} = 150\,\mathrm{GeV}$	89.6 ± 6.3
ZZ	384 ± 22
WZ	151.3 ± 9.4
Nonresonant bkg.	68 ± 17
Drell–Yan	70 ± 45
Other bkg.	14.7 ± 1.6
Total bkg.	688 ± 38
Data	694

Table 8.1: Signal predictions, post-fit background estimates, and observed numbers of events in the signal region. The combined statistical and systematic uncertainties are reported. Significant correlations exist between the various background yields, hence the sum of individual background uncertainties does not match the total background uncertainty.

are as predicted under the $\mu = 1$ hypothesis. The observed number of events in the ee ($\mu\mu$) channel is 286 (408), and the number of events expected from simulation is 300 ± 25 (388 ± 28). The per-channel expectation is estimated by dividing the post-fit $\ell\ell$ yield proportionally according to the pre-fit per-channel expectations from simulation. Several systematic uncertainties affect the two channels by differing amounts.

The total background estimates and the observed numbers of events in each $p_{\mathrm{T}}^{\mathrm{miss}}$ bin are listed in Table 8.2, for both a combined background-only fit to the signal region and the control regions, as well as for a fit to the control regions only. The latter results can be used in conjunction with the signal region bin correlation matrix presented in Fig. 8.11 to recast these results in the simplified likelihood framework (see Section 8.6).

Figure 8.1 shows the $p_{\mathrm{T}}^{\mathrm{miss}}$ distribution in the signal region (as well as the Drell–Yan sideband control region) after the full selection.

$p_{\mathrm{T}}^{\mathrm{miss}}$ bin (GeV)	Observed events	Total background prediction	
		SR+CR fit	CR-only fit
$100 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 125$	307	301 ± 33	259 ± 56
$125 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 150$	157	153 ± 14	147 ± 14
$150 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 175$	86	91.1 ± 6.2	88 ± 10
$175 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 200$	51	52.0 ± 3.6	50.3 ± 5.8
$200 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 250$	55	50.6 ± 2.8	49.8 ± 5.0
$250 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 300$	14	20.2 ± 1.3	19.8 ± 2.4
$300 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 350$	11	9.86 ± 0.74	9.7 ± 1.2
$350 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 400$	6	4.66 ± 0.37	4.55 ± 0.64
$400 \leq p_{\mathrm{T}}^{\mathrm{miss}} < 500$	6	3.84 ± 0.38	3.75 ± 0.60
$p_{\mathrm{T}}^{\mathrm{miss}} \geq 500$	1	1.88 ± 0.25	1.84 ± 0.38

Table 8.2: Expected event yields in each $p_{\mathrm{T}}^{\mathrm{miss}}$ bin for the sum of background processes in the signal region. The background yields and their corresponding uncertainties are obtained after performing a fit to data. Two sets of background yields are reported: one from a background-only fit to data in both the signal region and the control regions, and one from a fit to data in all control regions, but excluding data in the signal region. The observed numbers of events in each bin are also included.

No deviation from the Standard Model background expectation is found. In the following sections, exclusion limits on relevant model parameters are presented. These limits represent the threshold value of a given model parameter for which the 95% CL upper limit on signal strength (μ in Eqn. 7.2) crosses unity. Model parameters either above or below the threshold value are excluded, with the direction of exclusion dictated by the sign of the parameter's correlation with the signal cross section. The 95% CL upper limit on signal strength is computed as described in Section 7.8. In some cases, to facilitate comparison to results of other experimental searches, the 90% CL limit is used, which corresponds to finding the value of μ in Eqn. 7.4 for which $\mathrm{CL}_s = 1 - 0.9$.

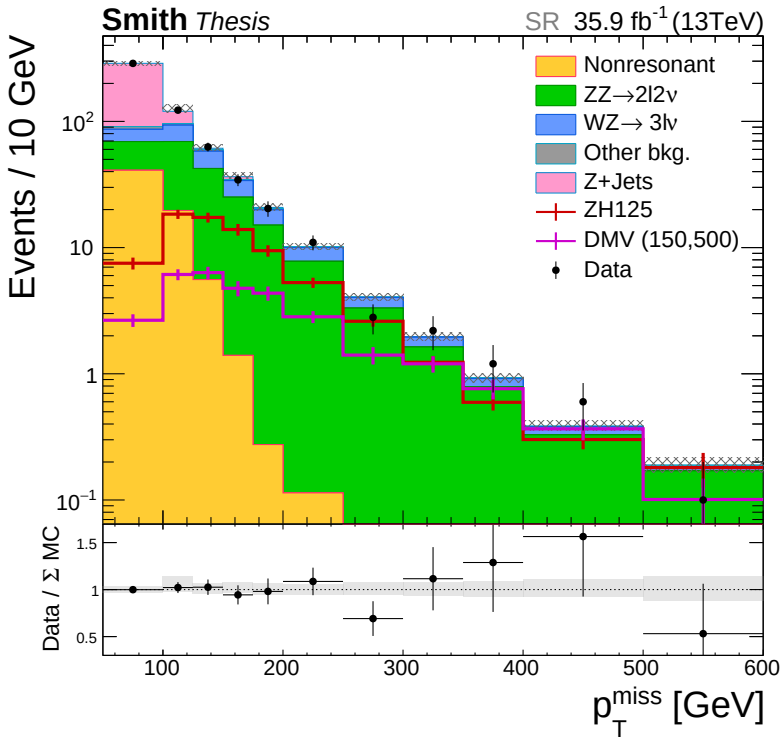

Figure 8.1: Post-fit distribution of the $p_{\mathrm{T}}^{\mathrm{miss}}$ in the signal region, after the full selection. The last bin also includes any events with $p_{\mathrm{T}}^{\mathrm{miss}} > 600\,\mathrm{GeV}$. The uncertainty band includes both statistical and systematic components. The ZH(inv.) signal normalization assumes SM production rates and the branching fraction $\mathcal{B}(\mathrm{H} \to \mathrm{inv.}) = 1$.

8.3 Particle dark matter interpretation

Figure 8.2 shows the 95% CL expected and observed exclusions in the $(m_{\mathrm{med}}, m_{\mathrm{DM}})$ plane for vector and axial vector mediated dark matter (DM) scenarios as described in Section 2.4.2, with couplings $g_q = 0.25$, $g_{\mathrm{DM}} = 1$ as recommended by the LHC dark matter working group [35] (see Section 5.3 for further details). The exclusions are found by establishing signal strength limits for DM simulation samples produced with mass parameters evenly spaced in a grid in the transformed coordinate system $(m_{\mathrm{med}}, 2m_{\mathrm{DM}}/m_{\mathrm{med}})$, and interpolating between these values via Delaunay triangulation [194]. The $\mu = 1$ contour is found, and then transformed back to the $(m_{\mathrm{med}}, m_{\mathrm{DM}})$

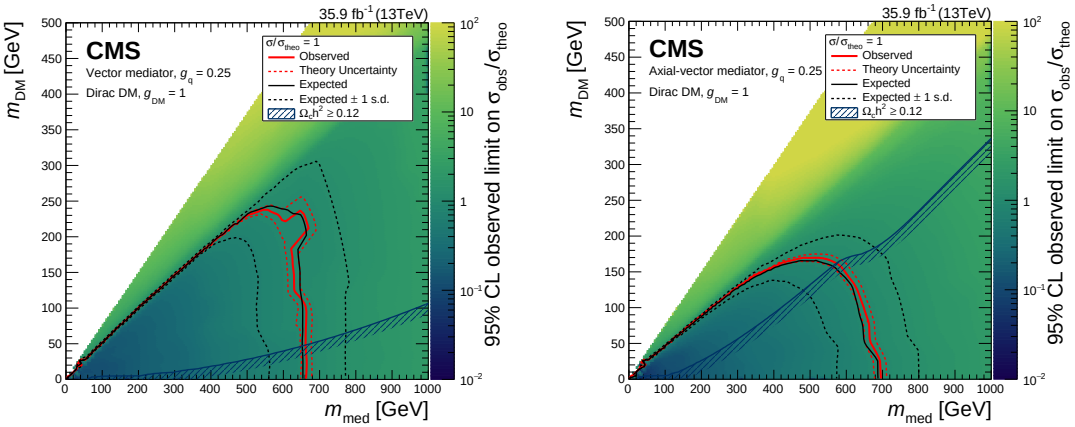

Figure 8.2: The 95% CL expected and observed limits on $\sigma_{\text{obs}}/\sigma_{\text{theo}}$ for the vector (left) and axial vector (right) mediators with $g_q = 0.25$ and $g_{\text{DM}} = 1$. Limits are not shown for far off-shell ($2m_{\text{DM}} > 1.5m_{\text{med}}$) regions of the parameter space. In addition, the thermal relic density constraint, assuming a single species of DM in the universe, is displayed according to Ref. [35].

coordinates. A limited subset of the DM simulation samples are produced with the full detector simulation as described in Section 5.6, while the remainder are generated by reweighting the reconstructed $p_{\text{T}}^{\text{miss}}$ spectrum of the nearest (in the $(m_{\text{med}}, m_{\text{DM}})$ plane) full simulation sample, with the weight being the ratio of the generated $p_{\text{T}}^{\text{miss}}$ in the target sample to that of the reference sample. The closure of this method is evaluated by evaluating target samples that were generated with the full simulation. The resulting limits are compatible with those found from the full simulation of these samples.

Figure 8.3 shows the 95% CL expected and observed limits on signal strength $\mu = \sigma_{\text{obs}}/\sigma_{\text{theo}}$ for couplings $g_q = g_{\text{DM}} = 1$ in the scalar and pseudoscalar mediated DM scenarios, as a function of m_{med} for a fixed DM mass of 1 GeV. As no signal strength limits are below unity, no region of $(m_{\text{DM}}, m_{\text{med}})$ parameter space is excluded by this result for these models.

The exclusions on m_{med} are computed at 90% CL as a function of the DM particle

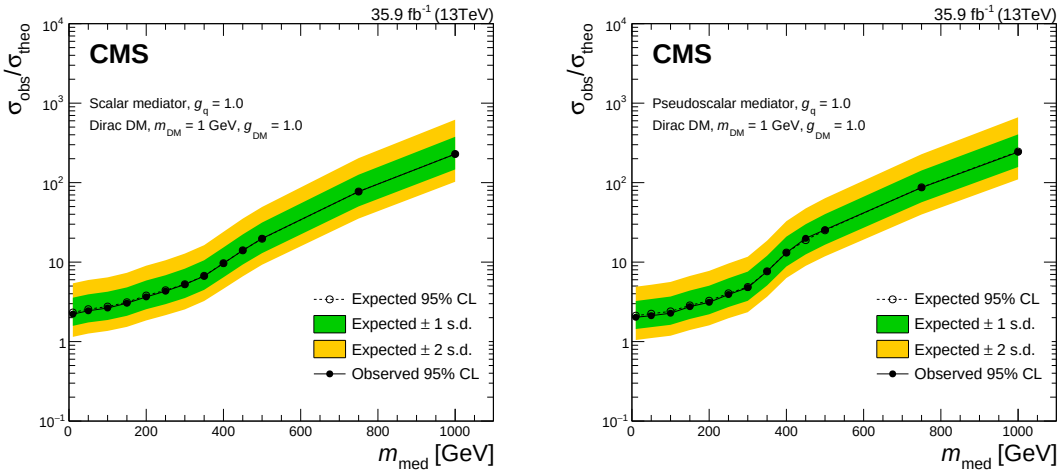

Figure 8.3: The 95% CL expected and observed limits on $\sigma_{\text{obs}}/\sigma_{\text{theo}}$ for the scalar (left) and pseudoscalar (right) mediated DM scenario with $g_q = g_{\text{DM}} = 1$. The limits are parameterized as a function of mediator mass m_{med} for a fixed dark matter mass $m_{\text{DM}} = 1\,\text{GeV}$.

mass, and translated according to Eqn. 2.5 to limits on the DM-nucleon scattering cross section, shown in Fig. 8.4. Both axial vector (spin-dependent) and vector (spin-independent) cases are considered. These exclusions are then compared to selected results from direct and indirect detection experiments (Section 3.2). The limits placed by this analysis are only applicable in the specific model described in Section 2.4.2, while the limits from direct and indirect detection experiments are independent of the assumed DM model. The limits placed by direct and indirect detection experiments assume that only one species of DM accounts for the observed relic density, whereas the limits placed by this analysis are independent of this assumption.

The exclusions presented in Figs. 8.2 and 8.3 are compared to the exclusions set by other CMS analysis targeting different final states that were produced concurrently with this result in Fig. 8.5. These exclusions are directly comparable as all analyses utilize the same simplified DM models with a common choice of coupling parameters. The absolute exclusions of these analyses, as well as their relative importance, will

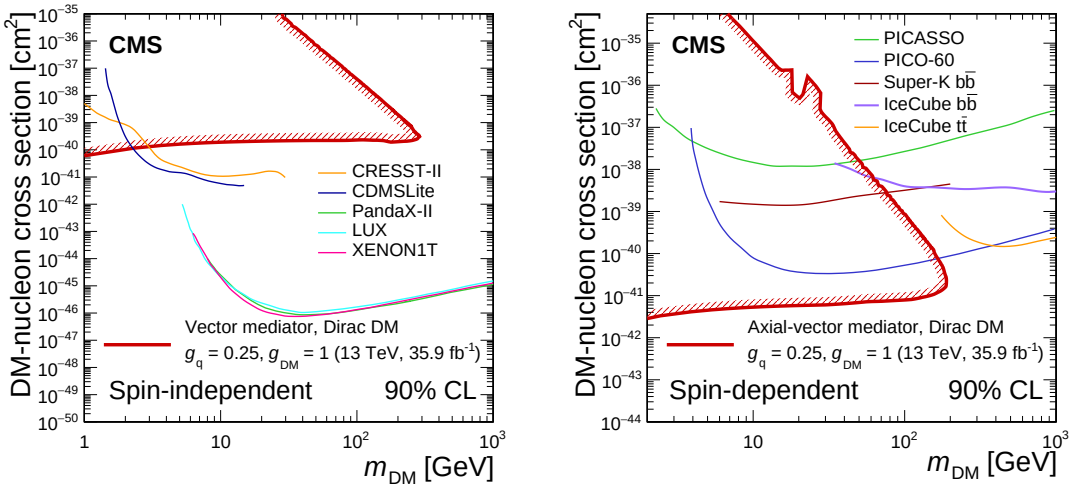

Figure 8.4: Observed 90% CL limits on the DM-nucleon scattering cross sections in both spin-independent (left) and spin-dependent (right) cases, assuming a mediator-quark coupling constant $g_q = 0.25$ and mediator-DM coupling constant $g_{\mathrm{DM}} = 1$. Limits from the CRESST-II [195], CDMSLite [196], PandaX-II [197], LUX [198], and XENON1T [199] experiments are shown for the spin-independent case (vector couplings). Limits from the PICASSO [200], PICO-60 [71], Super-Kamiokande [68], and IceCube [69, 70] experiments are shown for the spin-dependent case (axial vector couplings).

be altered by alternative choices of the DM model and/or coupling parameters. For example, with the choice of couplings given here, the relative sensitivities of the mono-Jet, mono-Photon, and mono-Z channels to spin-1 models is governed primarily by the relative probabilities of initial state jet, photon, and Z boson radiation. For the spin-0 models, the mono-Z sensitivity is much closer to that of mono-Jet, and both are subdominant to the $t\bar{t} + p_{\mathrm{T}}^{\mathrm{miss}}$ channel in some portions of the parameter space. For certain models involving coupling of a spin-0 mediator to the gauge sector, the mono-Z channel becomes more sensitive than the mono-Jet channel [62].

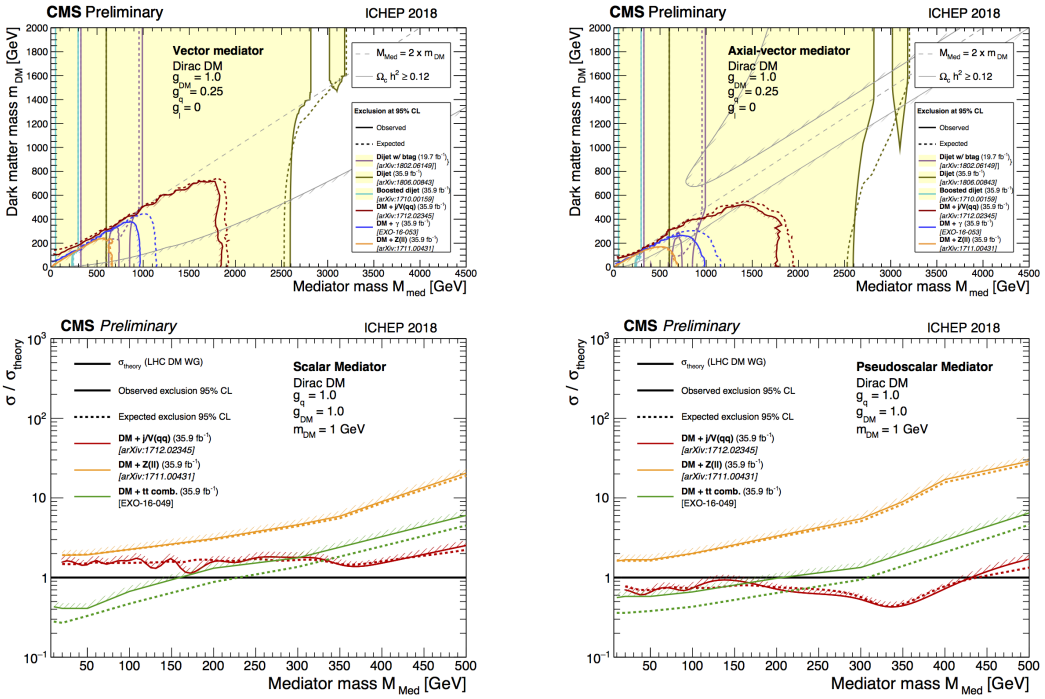

Figure 8.5: The 95% CL expected and observed exclusions for simplified dark matter models with vector (top left), axial vector (top right), scalar (bottom left), and pseudoscalar (bottom right) mediators, as found by di-jet searchs and various $p_{\mathrm{T}}^{\mathrm{miss}}$-based searches performed by the CMS collaboration. In all scenarios, $g_{\mathrm{DM}} = 1$, and for the vector and axial vector mediators, $g_q = 0.25$, while for the scalar and pseudoscalar mediators, $g_q = 1$. The absolute exclusion of the different searches, as well as their relative importance, will strongly depend on the chosen coupling and model scenario. Therefore, the exclusion regions and relic density contours shown in this plot are not applicable to other choices of coupling values or model.

8.4 Large extra dimension and unparticle interpretations

In the framework of the ADD model of large extra dimensions, limits are calculated depending on the number of extra dimensions n and the true Planck scale M_{D}. For each value of n, cross section limits are calculated as a function of M_{D}. By finding the intersection between the theory cross section line, calculated in the fiducial phase space of the graviton transverse momentum $p_{\mathrm{T}}^{G} > 50\,\mathrm{GeV}$, with the observed and

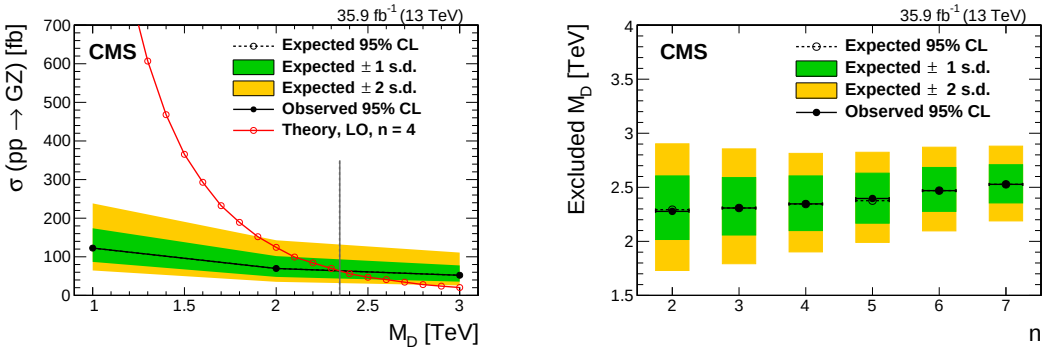

Figure 8.6: Expected and observed 95% CL cross section exclusion limits as a function of M_{D} for the example case $n = 4$ in the ADD scenario (left), and lower exclusion limits on M_{D} as a function of n (right). In both plots, the markers for the expected exclusion are obscured by the close overlap with those for the observed exclusion. The red solid line in the left plot shows the theoretical cross section for the case $n = 4$. Cross sections are calculated in the fiducial phase space of $p_{\mathrm{T}}^{G} > 50\,\mathrm{GeV}$. The vertical line in the left plot shows the projection onto the M_{D} axis of the intersection of the theory curve with the expected and observed exclusion limits.

expected excluded cross sections, and projecting that point onto the M_{D} axis, lower limits on M_{D} are set as a function of n, as shown in Fig. 8.6. In order to ensure the validity of the EFT, the ADD signal is truncated following Eqn. 2.7. Section 5.3 contains further details regarding the parameters used in the generation of the ADD signal simulation. The observed and expected exclusion of M_{D} ranges between 2.3 and 2.5 TeV for n between 2 and 7, at 95% CL. These exclusions are directly comparable to those presented in Fig. 3.3, where it is seen that results in the jet+$p_{\mathrm{T}}^{\mathrm{miss}}$ exclude a higher M_{D} for any given n than the result shown here. Nevertheless, this result can be used in a combination with those of other channels to improve the exclusion limits.

In the unparticle scenario, upper limits are set at 95% CL on the Wilson coefficient $\lambda/\Lambda_{\mathcal{U}}^{d_{\mathcal{U}}-1}$ of the unparticle-quark coupling operator of Eqn. 2.8, and are shown in Fig. 8.7 as a function of the scaling dimension $d_{\mathcal{U}}$. The excluded parameter space of this model depends strongly on the choice of operator, as discussed in Section 3.3.

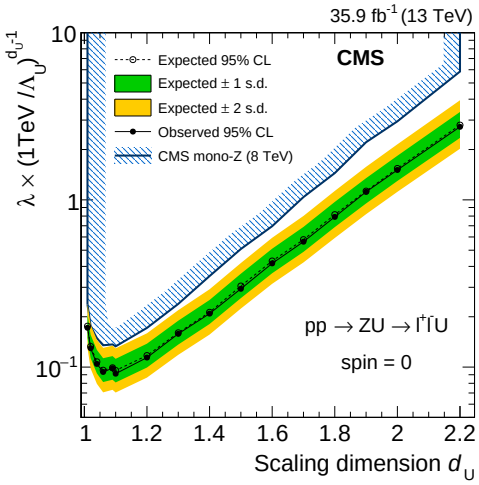

Figure 8.7: The 95% CL upper limits on the Wilson coefficient $\lambda/\Lambda_{\mathcal{U}}^{d_{\mathcal{U}}-1}$ of the unparticle-quark coupling operator. The results from an earlier CMS search in the same final state [57] are shown for comparison.

The only directly comparable results are the 8 TeV CMS searches in this final state [57] and in the jet+$p_{\mathrm{T}}^{\mathrm{miss}}$ final state [29]. This search excludes a factor of two larger parameter space than the 8 TeV result in this final state, and has approximately equal sensitivity as the 8 TeV jet+$p_{\mathrm{T}}^{\mathrm{miss}}$ result.

8.5 Invisible Higgs boson decay interpretation

The 95% CL expected and observed upper limits on the product of the production cross section and the branching fraction, $\sigma_{\mathrm{ZH}}\,\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ are shown as a function of the SM-like Higgs boson mass in Fig. 8.8. For $m_{\mathrm{H}} = 125\,\mathrm{GeV}$, the search can be interpreted as an upper limit on $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ assuming the SM production rate of a Higgs boson in association with a Z boson. Assuming the SM production rate, the 95% observed (expected) CL upper limit on $\mathcal{B}(\mathrm{H} \to \mathrm{inv.})$ is 0.45 (0.44). The $\mathrm{gg} \to \mathrm{Z}(\ell\ell)\mathrm{H}$ process is considered only for the 125 GeV mass point, and only when interpreting the result as a limit on branching fraction.

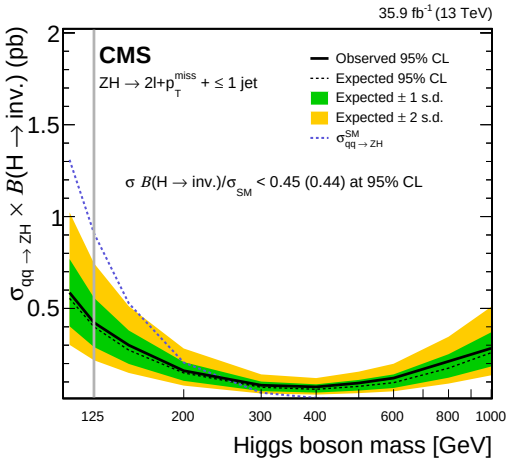

Figure 8.8: Expected and observed 95% CL upper limits on the product of the production cross section and the branching fraction, $\sigma_{qq\to\mathrm{ZH}}\,\mathcal{B}(\mathrm{H}\to\mathrm{inv.})$, as a function of the SM-like Higgs boson mass. The limits consider only quark-induced Higgs boson production. In addition, for the SM (125 GeV) Higgs boson, the limit on branching fraction assuming SM production rate (considering also gluon fusion) is presented. The vertical gray line indicates that the result at $m_{\mathrm{H}} = 125\,\mathrm{GeV}$ should not be read from the plot, as the gluon contribution is known for that point.

Following Section 2.6.2, the observed upper limit on $\mathcal{B}(\mathrm{H}\to\mathrm{inv.})$ for the SM Higgs boson is re-cast to a limit on the DM-nucleon cross section as a function of DM particle mass in the context of the Higgs-portal dark matter models presented in Eqn. 2.9. Fig. 8.9 shows the phase space of these models excluded by this result at the 90% CL, and compares it to the excluded phase space from various direct detection experiments.

The result of this analysis is combined[1] with that of two other CMS analyses targeting the SM Higgs invisible signal that were produced concurrently, by merging their respective likelihood functions into a single statistical model, and a combined observed (expected) upper limit of $\mathcal{B}(\mathrm{H}\to\mathrm{inv.}) < 0.24(0.18)$ is found [201]. The

[1] The result of the multivariate analysis presented in [4] is used for the combination, and is not presented here. This multivariate analysis gives a modest 5% relative improvement in the expected upper limit on $\mathcal{B}(\mathrm{H}\to\mathrm{inv.})$ relative to the $p_{\mathrm{T}}^{\mathrm{miss}}$-shape analysis presented here.

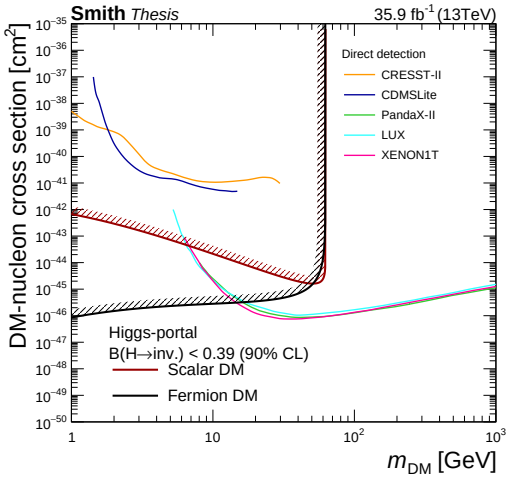

Figure 8.9: Observed 90% CL exclusions on the spin-independent DM-nucleon scattering cross sections for Higgs-portal dark matter models producing scalar or vector dark matter particles. The same spin-independent direct detection limits as in Fig. 8.4 are overlaid for comparison. The direct detection limits do not depend on the Higgs-mediator assumption.

relative contribution of this analysis is shown in Fig. 8.10. This result can be directly compared to that of Fig. 3.5, where an improvement in the expected limit in all channels is found. For this result, the observed limit in the VBF channel contains an approximately $1\,\sigma$ upward fluctuation.

8.6 Simplified likelihood

The results of this analysis are interpreted within a very limited subset of the new physics models for which this analysis may be sensitive. Typically, to interpret results of a given analysis in the context of a model not presented in the analysis, a model-independent limit is provided, where the 95% CL upper limit on the fiducial cross section of an arbitrary signal process is given for a specific choice of fiducial region. For an analysis such as this one, where the separation between signal and background is limited, a single fiducial region definition (i.e. a single $p_{\mathrm{T}}^{\mathrm{miss}}$ bin) cannot be chosen

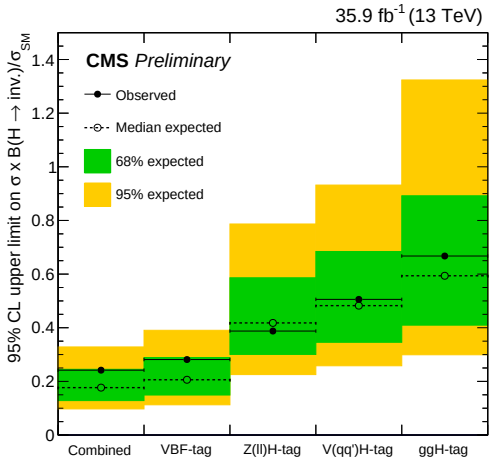

Figure 8.10: Observed and expected 95% CL limits on $\sigma\mathcal{B}(\mathrm{H} \to \mathrm{inv.})/\sigma_{\mathrm{SM}}$ for the VBF [201], Z($\ell\ell$)H [4], V(qq)H and gH [202] CMS H(inv.) analyses, as well as their combination.

optimally for all possible signal models, and in general will give poor results compared to those of a multi-bin shape analysis. The information necessary to perform the full multi-bin shape analysis (i.e. to reconstruct the likelihood function of Eqn. 7.2) is not suitable for publication, as the full effects of the detector acceptance and response to a given signal, as well as their uncertainties, must be considered, which requires technical knowledge beyond the scope of a re-interpretation. To allow re-interpretation of results such as this one, where the multi-bin analysis significantly enhances sensitivity to new physics signals, a simplified likelihood function is presented, following the prescription of Ref. [203].

The simplified likelihood function is constructed as

$$\mathcal{L} = \prod_i \mathcal{P}\left(N_{obs,i} | \mu N_{exp,i} + N_{bkg,i} + \theta_i\right) \cdot \exp\left(-\frac{1}{2}\vec{\theta}^T \left(\vec{\sigma}\mathbf{C}\vec{\sigma}^T\right)^{-1} \vec{\theta}\right),$$

where $N_{obs,i}$ and $N_{bkg,i}$ are the data and background yields in $p_{\mathrm{T}}^{\mathrm{miss}}$ bin i, $N_{exp,i}$ is the predicted signal yield for the alternative signal model in bin i, $\vec{\theta}$ is a set of nuisance parameters encoding deviations from the background expectations in each bin, which

are constrained by a prior probability density function that is approximated by a multivariate normal distribution formed from a CR-only fit result, i.e. the vector of per-bin background yield uncertainties $\vec{\sigma}$ and the correlation matrix of those uncertainties $\mathbf{C}$. This simplified likelihood function is maximized with respect to $\{\mu, \theta_i\}$, and the usual machinery of Section 7.8 is employed to find signal strength limits for an alternative signal hypothesis. By utilizing a CR-only fit result to form the prior probability density function of the background expectation, the likelihood can be maximized with respect to a new signal hypothesis without double-fitting of the observation in the signal region. The observed data counts and the total background estimates in the signal region from the control-region-only fit are shown in Table 8.2. Figure 8.11 shows the correlations between the estimated background yields in $p_{\mathrm{T}}^{\mathrm{miss}}$ bins in the signal region, for the control-region-only fit to the observed data.

To utilize the simplified likelihood method, a prediction of the reconstructed event yields in each $p_{\mathrm{T}}^{\mathrm{miss}}$ bin is required. This is best obtained by using a detector simulation program such as DELPHES [204], however a reasonable prediction can be obtained by: applying a generator-level selection that parallels the reconstruction-level selection described in Section 7.4, omitting tau lepton and b jet vetoes; smearing the $p_{\mathrm{T}}^{\mathrm{miss}}$ with a Gaussian kernel of 24 GeV width; and scaling by a reconstruction efficiency of 0.70.

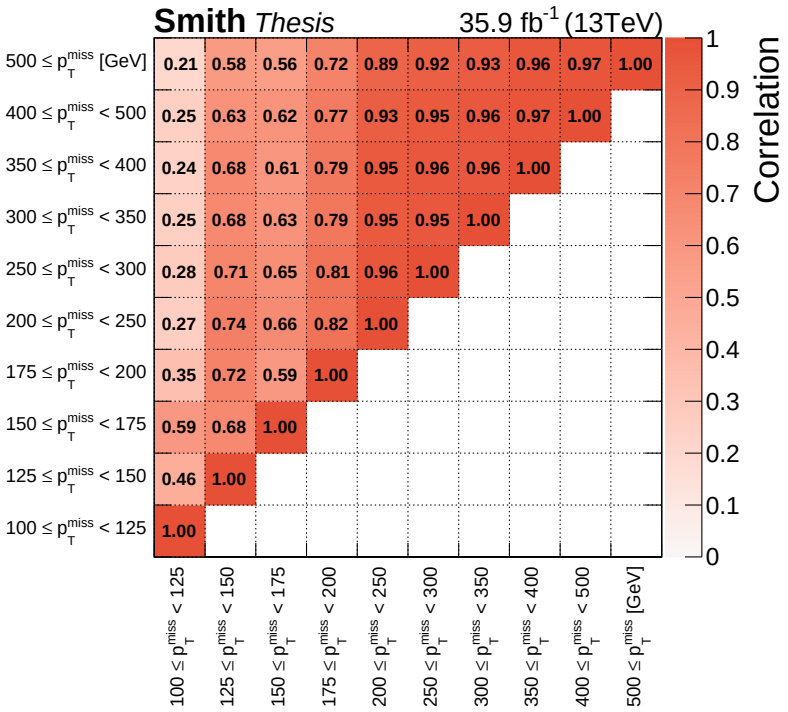

Figure 8.11: Correlations between the estimated background yields in the signal region $p_{\mathrm{T}}^{\mathrm{miss}}$ bins. The correlations are obtained after performing a combined fit to data in all control regions, but excluding data in the signal region. Since the correlation matrix is symmetric by construction, the part below the diagonal is not shown.

Chapter 9

Conclusions

9.1 Summary

A search for new physics in events with a leptonically decaying Z boson and a large transverse momentum imbalance has been presented. This search is based on a data set of proton-proton collisions collected with the CMS experiment in 2016, corresponding to an integrated luminosity of $35.9\,\mathrm{fb}^{-1}$ at $\sqrt{s} = 13\,\mathrm{TeV}$. No evidence for physics beyond the Standard Model is found. This result extends the excluded parameter space for simplified dark matter models with spin-1 mediators and an unparticle model by over a factor of two, compared to a previous search performed in this final state with the CMS detector [59]. This is due to the increase in available integrated luminosity and improvements in the analysis technique. For example, an observed (expected) lower exclusion limit on $m_{\mathrm{med}} < 690\,(690)\,\mathrm{GeV}$ was set on dark matter production in the simplified model with an axial vector mediator for $m_{\mathrm{DM}} = 1\,\mathrm{GeV}$, where the corresponding limit in the previous search was $300\,(270)\,\mathrm{GeV}$, for the particular choice of couplings as described in Section 8.3. Results for dark matter production via spin-0 mediators in the simplified model interpretation, as well as graviton emis-

sion in a model with large extra dimensions, are presented in this final state for the first time. An observed (expected) upper limit of 45% (44%) has been set on the branching ratio of the Standard Model Higgs boson to invisible final states at the 95% confidence level. The result presented in this analysis includes the necessary materials to re-interpret the observed data in the context of alternative models of new physics with significantly improved fidelity over traditional model-independent limits via the simplified likelihood method, a first for this final state.

A search for new physics in this final state with a dataset corresponding to a similar integrated luminosity ($36.1\,\mathrm{fb}^{-1}$) was performed by the ATLAS collaboration, where an observed (expected) upper limit of $\mathcal{B}(\mathrm{H} \to \mathrm{inv.}) < 67\%\,(39\%)$ was set, and an observed (expected) lower exclusion limit on $m_{\mathrm{med}} < 560\,(620)\,\mathrm{GeV}$ was set on dark matter production in the simplified model with an axial vector mediator for $m_{\mathrm{DM}} = 1\,\mathrm{GeV}$ and with the same coupling parameters as in this analysis [205].

9.2 Outlook

Run 2 of the LHC is coming to a close this year, and with it an integrated luminosity in excess of $150\,\mathrm{fb}^{-1}$ of proton-proton collisions at $\sqrt{s} = 13\,\mathrm{TeV}$ will become available for analysis. This analysis will benefit from the at least 4-fold increase in luminosity, as the control region yields will become statistically significant, and hence will better constrain the diboson background that is a major limitation on the sensitivity to new physics in this final state. The theoretical predictions for all Standard Model backgrounds are continuously improving with time, and with them the ability to utilize more aggressive selections and assign reduced systematic uncertainties will also improve the sensitivity to new physics in this analysis. Additional improvements in the simulation of the CMS detector and reconstruction algorithms can be expected

in the coming years, and this will likely improve the modeling of backgrounds, in particular the Drell–Yan plus mis-reconstructed $p_{\mathrm{T}}^{\mathrm{miss}}$ background. Future iterations of this analysis may be optimized to target alternative DM models involving mediator coupling to the gauge sector, which can enhance the cross section for $\mathrm{Z} + p_{\mathrm{T}}^{\mathrm{miss}}$ production relative to that of other initial state radiation based searches for dark matter [62].

The operation of the High-Luminosity Large Hadron Collider (HL-LHC) is expected to begin in 2025. The HL-LHC will collect data with an instantaneous luminosity 5 times that of the LHC, and will run for longer periods of time, such that it may collect $3000\,\mathrm{fb}^{-1}$ of integrated luminosity over its planned 10 year lifetime. This amount of data is expected to exclude significantly more parameter space of the models probed in this analysis, and the discovery reach (i.e. the region of parameter space in which a $5\,\sigma$ significance for a hypothetical excess could be obtained) will be enhanced.

Searches for new physics at hadron colliders benefit most from increases in center of mass energy rather than luminosity, since the ability of hadron collider experiments to probe new physics at a given energy scale is limited by the rate of production of parton-parton collisions at that energy scale. This rate is of course dependent on the final state, but can be approximated to sufficient accuracy to make qualitative statements about the performance of a hadron collider as a function of center of mass energy and integrated luminosity. This is accomplished via the parton luminosity distribution, which describes the approximate differential cross section for a parton-parton interaction $i + j \to X$ as a function of $\sqrt{\hat{s}}$ (Eqn. 2.3) up to a dimensionless factor $\hat{s}\sigma_{ij\to X}(\hat{s})$ that is typically of order unity [206]. Parton luminosities for gluon–gluon and up quark–down antiquark interactions are shown in Fig. 9.1. The $u\bar{d}$ parton luminosity is similar to that of $u\bar{u}$ and other valence quark–sea antiquark

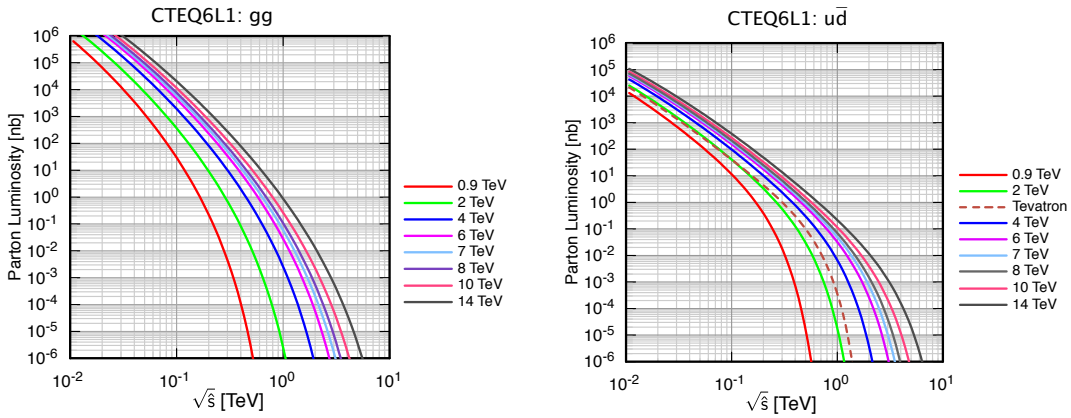

Figure 9.1: Parton luminosities for gg and $u\bar{d}$ interactions as a function of parton-parton center of mass energy $\sqrt{\hat{s}}$, for various proton-proton center of mass energies $\sqrt{s}$. Reproduced from [206], Figures 1 and 2.

combinations (Section 5.2), which, for example, approximate the cross section of dark matter mediator production. In both cases, one observes an inflection point in the graph, beyond which the parton luminosity begins to fall at such a rate that the relative gain in $\sqrt{\hat{s}}$ sensitivity for a given increase in integrated luminosity is marginal. The best way to probe higher energy scales is to increase $\sqrt{s}$, and recalling Eqn. 4.1, this requires an increase the magnetic field of the bending magnets and/or the radius of the collider.

Bibliography

[1] E. Rutherford. The scattering of alpha and beta particles by matter and the structure of the atom. *Phil. Mag. Ser.6*, 21:669–688, 1911. doi:10.1080/14786440508637080.

[2] G. Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B*, 716:1, 2012. arXiv:1207.7214.

[3] S. Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B*, 716:30, 2012. arXiv:1207.7235.

[4] CMS Collaboration. Search for new physics in events with a leptonically decaying Z boson and a large transverse momentum imbalance in proton-proton collisions at $\sqrt{s} = 13$ TeV. *Eur. Phys. J. C*, 78:291, 2018. arXiv:1711.00431.

[5] F. Halzen and Alan D. Martin. *Quarks and leptons: An introductory course in modern particle physics*. Wiley, New York, USA, 1984.

[6] David J. Griffiths. *Introduction to elementary particles*. Wiley, New York, USA, 1987.

[7] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory*. Addison-Wesley, Reading, USA, 1995.

[8] M. Srednicki. *Quantum field theory*. Cambridge University Press, 2007.

[9] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016. doi:10.1088/1674-1137/40/10/100001.

[10] V. C. Rubin, N. Thonnard, and W. K. Ford, Jr. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605

(R = 4kpc) to UGC 2885 (R = 122 kpc). *Astrophys. J.*, 238:471, 1980. doi:10.1086/158003.

[11] K. G. Begeman, A. H. Broeils, and R. H. Sanders. Extended rotation curves of spiral galaxies: Dark haloes and modified dynamics. *Mon. Not. Roy. Astron. Soc.*, 249:523, 1991. doi:10.1093/mnras/249.3.523.

[12] Gianfranco Bertone, Dan Hooper, and Joseph Silk. Particle dark matter: Evidence, candidates and constraints. *Phys. Rept.*, 405:279–390, 2005. arXiv:hep-ph/0404175.

[13] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helv. Phys. Acta*, 6:110–127, 1933. doi:10.1007/s10714-008-0707-4.

[14] A. Einstein. Die Feldgleichungen der Gravitation. *Königlich Preußische Akademie der Wissenshaften*, page 844, 1915. Einstein papers online.

[15] A. Friedman. Über die krümmung des raumes. *Zeitschrift für Physik*, 10(1):377–386, Dec 1922. doi:10.1007/BF01332580.

[16] George F. Smoot. COBE observations and results. *AIP Conf. Proc.*, 476(1):1–10, 1999. arXiv:astro-ph/9902027.

[17] P. Sarkar, J. Yadav, B. Pandey, and S. Bharadwaj. The scale of homogeneity of the galaxy distribution in SDSS DR6. *Monthly Notices of the Royal Astronomical Society: Letters*, 399:L128–L131, 2009. arXiv:1610.07922.

[18] G. Hinshaw et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Cosmological parameter results. *Astrophys. J. Suppl.*, 208:19, 2013. arXiv:1212.5226.

[19] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594:A13, 2016. arXiv:1502.01589.

[20] Simon Knapen, Tongyan Lin, and Kathryn M. Zurek. Light Dark Matter: Models and Constraints. *Phys. Rev.*, D96(11):115021, 2017. arXiv:1709.07882.

[21] Maria Beltran, Dan Hooper, Edward W. Kolb, Zosia A. C. Krusberg, and Tim M. P. Tait. Maverick dark matter at colliders. *JHEP*, 09:037, 2010. arXiv:1002.4137.

[22] Jessica Goodman, Masahiro Ibe, Arvind Rajaraman, William Shepherd, Tim M. P. Tait, and Hai-Bo Yu. Constraints on Light Majorana dark Matter from Colliders. *Phys. Lett. B*, 695:185–188, 2011. arXiv:1005.1286.

[23] Yang Bai, Patrick J. Fox, and Roni Harnik. The Tevatron at the Frontier of Dark Matter Direct Detection. *JHEP*, 12:048, 2010. arXiv:1005.3797.

[24] Jessica Goodman, Masahiro Ibe, Arvind Rajaraman, William Shepherd, Tim M. P. Tait, and Hai-Bo Yu. Constraints on Dark Matter from Colliders. *Phys. Rev. D*, 82:116010, 2010. arXiv:1008.1783.

[25] Enrico Fermi. Tentativo di una teoria dell'emissione dei raggi beta. *Ric. Sci.*, 4:491–495, 1933. English translation: doi:10.1119/1.1974382.

[26] Steven Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[27] G. Arnison et al. Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at s**(1/2) = 540 GeV. *Phys. Lett. B*, 122:103–116, 1983. doi:10.1016/0370-2693(83)91177-2.

[28] M. Banner et al. 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*, 122:476–485, 1983. doi:10.1016/0370-2693(83)91605-2.

[29] CMS Collaboration. Search for dark matter, extra dimensions, and unparticles in monojet events in proton-proton collisions at $\sqrt{s} = 8$ TeV. *Eur. Phys. J. C*, 75(5):235, 2015. arXiv:1408.3583.

[30] ATLAS Collaboration. Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp collisions at $\sqrt{s}$ =8 TeV with the ATLAS detector. *Eur. Phys. J. C*, 75(7):299, 2015. [Erratum: Eur. Phys. J.C75,no.9,408(2015)] arXiv:1502.01518.

[31] Giorgio Busoni, Andrea De Simone, Enrico Morgante, and Antonio Riotto. On the Validity of the Effective Field Theory for Dark Matter Searches at the LHC. *Phys. Lett. B*, 728:412–421, 2014. arXiv:1307.2253.

[32] Christoph Englert and Michael Spannowsky. Effective Theories and Measurements at Colliders. *Phys. Lett. B*, 740:8–15, 2015. arXiv:1408.5147.

[33] Daniel Abercrombie et al. Dark Matter Benchmark Models for Early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum. 2015. arXiv:1507.00966.

[34] Jalal Abdallah et al. Simplified Models for Dark Matter Searches at the LHC. *Phys. Dark Univ.*, 9-10:8–23, 2015. arXiv:1506.03116.

[35] Giorgio Busoni et al. Recommendations on presenting LHC searches for missing transverse energy signals using simplified s-channel models of dark matter. 2016. arXiv:1603.04156.

[36] Nima Arkani-Hamed, Savas Dimopoulos, and G. R. Dvali. The hierarchy problem and new dimensions at a millimeter. *Phys. Lett. B*, 429:263, 1998. arXiv:hep-ph/9803315.

[37] Nima Arkani-Hamed, Savas Dimopoulos, and G. R. Dvali. Phenomenology, astrophysics and cosmology of theories with submillimeter dimensions and TeV scale quantum gravity. *Phys. Rev. D*, 59:086004, 1999. arXiv:hep-ph/9807344.

[38] Tao Han, Joseph D. Lykken, and Ren-Jie Zhang. On Kaluza-Klein states from large extra dimensions. *Phys. Rev. D*, 59:105006, 1999. arXiv:hep-ph/9811350.

[39] Howard Georgi. Unparticle physics. *Phys. Rev. Lett.*, 98:221601, 2007. arXiv:hep-ph/0703260.

[40] Gian F. Giudice, Riccardo Rattazzi, and James D. Wells. Quantum gravity and extra dimensions at high-energy colliders. *Nucl. Phys. B*, 544:3–38, 1999. arXiv:hep-ph/9811291.

[41] S. Ask, I. V. Akin, L. Benucci, A. De Roeck, M. Goebel, and J. Haller. Real Emission and Virtual Exchange of Gravitons and Unparticles in Pythia8. *Comput. Phys. Commun.*, 181:1593–1604, 2010. arXiv:0912.4233.

[42] Tom Banks and A. Zaks. On the phase structure of vector-like gauge theories with massless fermions. *Nucl. Phys. B*, 196:189, 1982. doi:10.1016/0550-3213(82)90035-9.

[43] Kingman Cheung, Wai-Yee Keung, and Tzu-Chiang Yuan. Collider Phenomenology of Unparticle Physics. *Phys. Rev.*, D76:055003, 2007. arXiv:0706.3155.

[44] ATLAS and CMS Collaborations. Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV. *JHEP*, 08:045, 2016. arXiv:1606.02266.

[45] Diptimoy Ghosh, Rohini Godbole, Monoranjan Guchait, Kirtimaan Mohan, and Dipan Sengupta. Looking for an invisible Higgs signal at the LHC. *Phys. Lett. B*, 725:344, 2013. arXiv:1211.7015.

[46] Stephen P. Martin and James D. Wells. Motivation and detectability of an invisibly decaying Higgs boson at the Fermilab Tevatron. *Phys. Rev. D*, 60:035006, 1999. arXiv:hep-ph/9903259.

[47] Yang Bai, Patrick Draper, and Jessie Shelton. Measuring the invisible Higgs width at the 7 and 8 TeV LHC. *JHEP*, 07:192, 2012. arXiv:1112.4496.

[48] Abdelhak Djouadi, Oleg Lebedev, Yann Mambrini, and Jeremie Quevillon. Implications of LHC searches for Higgs–portal dark matter. *Phys. Lett. B*, 709:65, 2012. arXiv:1112.3299.

[49] Seungwon Baek, P. Ko, Wan-Il Park, and Eibun Senaha. Higgs portal vector dark matter: revisited. *JHEP*, 5:36, 2013. arXiv:1212.2131.

[50] Martin Hoferichter, Philipp Klos, Javier Menéndez, and Achim Schwenk. Improved limits for Higgs-portal dark matter from LHC searches. *Phys. Rev. Lett.*, 119(18):181803, 2017. arXiv:1708.02245.

[51] Gian F. Giudice, Riccardo Rattazzi, and James D. Wells. Graviscalars from higher dimensional metrics and curvature Higgs mixing. *Nucl. Phys. B*, 595:250, 2001. arXiv:hep-ph/0002178.

[52] M. Battaglia, D. Dominici, J. F. Gunion, and J. D. Wells. The Invisible Higgs decay width in the add model at the LHC. In *Physics at TeV colliders. Proceedings, Workshop, Les Houches, France, May 26-June 3, 2003*, 2004. arXiv:hep-ph/0402062.

[53] Antonio Delgado, Jose R. Espinosa, and Mariano Quiros. Unparticles Higgs Interplay. *JHEP*, 10:094, 2007. arXiv:0707.4309.

[54] D. de Florian et al. Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector. 2016. arXiv:1610.07922.

[55] J. Abdallah et al. Searches for invisibly decaying Higgs bosons with the DELPHI detector at LEP. *Eur. Phys. J.*, C32:475–492, 2004. arXiv:hep-ex/0401022.

[56] C. Principato. Search for an Invisible Decaying Higgs Boson in Dilepton Events at CDF. *Nucl. Part. Phys. Proc.*, 273-275:2476–2478, 2016. doi:10.1016/j.nuclphysbps.2015.09.430.

[57] CMS Collaboration. Search for dark matter and unparticles produced in association with a Z boson in proton-proton collisions at $\sqrt{s} = 8$ TeV. *Phys. Rev.*, D93(5):052011, 2016. arXiv:1511.09375.

[58] ATLAS Collaboration. Search for dark matter in events with a Z boson and missing transverse momentum in pp collisions at $\sqrt{s}$=8 TeV with the ATLAS detector. *Phys. Rev.*, D90(1):012004, 2014. arXiv:1404.0051.

[59] CMS Collaboration. Search for dark matter and unparticles in events with a Z boson and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV. *JHEP*, 03:061, 2017. [Erratum: JHEP09,106(2017)], arXiv:1701.02042.

[60] CMS Collaboration. Search for ZZ resonances in the $2\ell 2\nu$ final state in proton-proton collisions at 13 TeV. *JHEP*, 03:003, 2018. arXiv:1711.04370.

[61] CMS Collaboration. Search for a new scalar resonance decaying to a pair of Z bosons in proton-proton collisions at $\sqrt{s} = 13$ TeV. 2018. arXiv:1804.01939.

[62] Martin Bauer, Ulrich Haisch, and Felix Kahlhoefer. Simplified dark matter models with two Higgs doublets: I. Pseudoscalar mediators. *JHEP*, 05:138, 2017. arXiv:1701.07427.

[63] CMS Collaboration. Search for new phenomena in monophoton final states in proton-proton collisions at $\sqrt{s} = 8$ TeV. *Phys. Lett.*, B755:102–124, 2016. arXiv:1410.8812.

[64] ATLAS Collaboration. Search for new phenomena in events with a photon and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *Phys. Rev.*, D91(1):012008, 2015. [Erratum: Phys. Rev.D92,no.5,059903(2015)], arXiv:1411.1559.

[65] CMS Collaboration. Search for dark matter produced with an energetic jet or a hadronically decaying W or Z boson at $\sqrt{s} = 13$ TeV. *JHEP*, 07:014, 2017. arXiv:1703.01651.

[66] Search for new physics in a boosted hadronic monotop final state using 12.9 fb^{-1} of $\sqrt{s} = 13$ TeV data. Technical Report CMS-PAS-EXO-16-040, CERN, Geneva, 2016. cds:CMS-PAS-EXO-16-040.

[67] Dark Matter Summary Plots from CMS for ICHEP 2016. Aug 2016. cds:CMS-DP-2016-057.

[68] K. Choi et al. Search for neutrinos from annihilation of captured low-mass dark matter particles in the Sun by Super-Kamiokande. *Phys. Rev. Lett.*, 114:141301, 2015. arXiv:1503.04858.

[69] M. G. Aartsen et al. Search for annihilating dark matter in the Sun with 3 years of IceCube data. *Eur. Phys. J. C*, 77:146, 2017. arXiv:1612.05949.

[70] M. G. Aartsen et al. Improved limits on dark matter annihilation in the Sun with the 79-string IceCube detector and implications for supersymmetry. *JCAP*, 04:022, 2016. arXiv:1601.00653.

[71] C. Amole et al. Dark matter search results from the PICO-60 C_3F_8 bubble chamber. *Phys. Rev. Lett.*, 118:251301, 2017. arXiv:1702.07666.

[72] R. Bernabei et al. New results from DAMA/LIBRA. *Eur. Phys. J.*, C67:39–49, 2010. arXiv:1002.1028.

[73] R. Agnese et al. Silicon Detector Dark Matter Results from the Final Exposure of CDMS II. *Phys. Rev. Lett.*, 111(25):251301, 2013. arXiv:1304.4279.

[74] G. Taylor. Single photon and multiphoton production in e+ e- collisions at a center-of-mass energy of 188.6-GeV. In *Proceedings, International Europhysics Conference on High energy physics (EPS-HEP 1999): Tampere, Finland, July 15-21, 1999*, 1999. cds:CERN-OPEN-99-268.

[75] G. Abbiendi et al. Photonic events with missing energy in e+ e- collisions at S**(1/2) = 189-GeV. *Eur. Phys. J.*, C18:253–272, 2000. arXiv:hep-ex/0005002.

[76] P. Achard et al. Single photon and multiphoton events with missing energy in e^+e^- collisions at LEP. *Phys. Lett.*, B587:16–32, 2004. arXiv:hep-ex/0402002.

[77] V. M. Abazov et al. Search for large extra dimensions in the monojet + missing E_T channel at DØ. *Phys. Rev. Lett.*, 90:251802, 2003. arXiv:hep-ex/0302014.

[78] D. Acosta et al. Search for Kaluza-Klein graviton emission in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$-TeV using the missing energy signature. *Phys. Rev. Lett.*, 92:121802, 2004. arXiv:hep-ex/0309051.

[79] A. Abulencia et al. Search for Large Extra Dimensions in the Production of Jets and Missing Transverse Energy in p anti-p Collisions at s**(1/2) = 1.96 TeV. *Phys. Rev. Lett.*, 97:171802, 2006. arXiv:hep-ex/0605101.

[80] T. M. Aliev, Alan S. Cornell, and Naveen Gaur. B —¿ K(K*) missing energy in Unparticle physics. *JHEP*, 07:072, 2007. arXiv:0705.4542.

[81] Bernard Aubert et al. Search for $B \to K^* \nu \bar{\nu}$ decays. *Phys. Rev.*, D78:072007, 2008. arXiv:0808.1338.

[82] J. P. Lees et al. Search for $B \to K^{(*)} \nu \overline{\nu}$ and invisible quarkonium decays. *Phys. Rev.*, D87(11):112005, 2013. arXiv:1303.7465.

[83] Yi Liao. Bounds on unparticles couplings to electrons: From electron g-2 to positronium decays. *Phys. Rev.*, D76:056006, 2007. arXiv:0705.0837.

[84] Shao-Long Chen, Xiao-Gang He, and Ho-Chin Tsai. Constraints on unparticle interactions from invisible decays of Z, quarkonia and neutrinos. *JHEP*, 11:010, 2007. arXiv:0707.0187.

[85] Robert E. Shrock and Mahiko Suzuki. Invisible Decays of Higgs Bosons. *Phys. Lett.*, 110B:250, 1982. doi:10.1016/0370-2693(82)91247-3.

[86] A. Lopez-Fernandez, J. C. Romao, F. de Campos, and J. W. F. Valle. Model independent Higgs boson mass limits at LEP. *Phys. Lett. B*, 312:240–246, 1993. arXiv:hep-ph/9304255.

[87] Searches for invisible Higgs bosons: Preliminary combined results using LEP data collected at energies up to 209-GeV. In *Lepton and photon interactions at high energies. Proceedings, 20th International Symposium, LP 2001, Rome, Italy, July 23-28, 2001*, 2001. arXiv:hep-ex/0107032.

[88] G. Abbiendi et al. Search for invisibly decaying Higgs bosons with large decay width using the OPAL detector at LEP. *Eur. Phys. J.*, C49:457–472, 2007. arXiv:hep-ex/0610056.

[89] R W Assmann. LEP Operation and Performance with Electron-Positron Collisions at 209 GeV. 2001. cds:CERN-SL-2001-003.

[90] ATLAS Collaboration. Constraints on new phenomena via Higgs boson couplings and invisible decays with the ATLAS detector. *JHEP*, 11:206, 2015. arXiv:1509.00672.

[91] CMS Collaboration. Searches for invisible decays of the Higgs boson in pp collisions at $\sqrt{s} = 7$, 8, and 13 TeV. *JHEP*, 02:135, 2017. arXiv:1610.09218.

[92] Lyndon Evans and Philip Bryant. LHC Machine. *JINST*, 3:S08001, 2008. doi:10.1088/1748-0221/3/08/S08001.

[93] T. Boutboul, S. Le Naour, D. Leroy, L. Oberli, and V. Previtali. Critical Current Density in Superconducting Nb-Ti trands in the 100 mT to 11 T Applied Field Range. *IEEE Trans. Appl. Supercond.*, 16(2):1184–1187, 2006. doi:10.1109/TASC.2006.870777.

[94] A. Verweij et al. Retraining of the 1232 Main Dipole Magnets in the LHC. *IEEE Trans. Appl. Supercond.*, 26(4):4000705, 2016.

[95] G. Aad et al. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008. doi:10.1088/1748-0221/3/08/S08003.

[96] A. Augusto Alves, Jr. et al. The LHCb Detector at the LHC. *JINST*, 3:S08005, 2008. doi:10.1088/1748-0221/3/08/S08005.

[97] K. Aamodt et al. The ALICE experiment at the CERN LHC. *JINST*, 3:S08002, 2008. doi:10.1088/1748-0221/3/08/S08002.

[98] S. Chatrchyan et al. The CMS Experiment at the CERN LHC. *JINST*, 3:S08004, 2008. doi:10.1088/1748-0221/3/08/S08004.

[99] P. Baudrenghien, T. Bohl, T. Linnecar, E. Shaposhnikova, and J. Tuckmantel. Nominal longitudinal parameters for the LHC beam in the CERN SPS. *Conf. Proc.*, C030512:3050, 2003. cds:CERN-LHC-PROJECT-REPORT-652.

[100] Oliver S. Bruning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole, and P. Proudlock. LHC Design Report Vol.1: The LHC Main Ring. 2004. cds:CERN-2004-003-V1.

[101] Roderik Bruce et al. LHC Run 2: Results and challenges. In *Proceedings, 57th ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (HB2016): Malmö, Sweden, July 3-8, 2016*, page MOAM5P50, 2016. doi:10.18429/JACoW-HB2016-MOAM5P50.

[102] ATLAS Collaboration. Measurement of the Inelastic Proton-Proton Cross Section at $\sqrt{s} = 13$ TeV with the ATLAS Detector at the LHC. *Phys. Rev. Lett.*, 117(18):182002, 2016. arXiv:1606.02625.

[103] S Chatrchyan et al. Commissioning of the CMS Experiment and the Cosmic Run at Four Tesla. *JINST*, 5:T03001, 2010. arXiv:0911.4845.

[104] V. I. Klyukhin et al. The CMS Magnetic Field Map Performance. *IEEE Trans. Appl. Supercond.*, 20(3):152–155, 2010. arXiv:1110.0607.

[105] S Chatrchyan et al. Precise Mapping of the Magnetic Field in the CMS Barrel Yoke using Cosmic Rays. *JINST*, 5:T03021, 2010. arXiv:0910.5530.

[106] P. Adzic et al. Energy resolution of the barrel of the CMS electromagnetic calorimeter. *JINST*, 2:P04004, 2007. doi:10.1088/1748-0221/2/04/P04004.

[107] P. Cushman, A. Heering, J. Nelson, Charles Timmermans, S. R. Dugad, S. Katta, and S. Tonwar. Multi-pixel hybrid photodiode tubes for the CMS hadron calorimeter. *Nucl. Instrum. Meth.*, A387:107–112, 1997. doi:10.1016/S0168-9002(96)00972-2.

[108] Victor Daniel Elvira. Measurement of the Pion Energy Response and Resolution in the CMS HCAL Test Beam 2002 Experiment. 2004. cds:CMS-NOTE-2004-020.

[109] S. Dasu et al. CMS. The TriDAS project. Technical design report, vol. 1: The trigger systems. 2000. cds:CERN-LHCC-2000-038.

[110] A. Tapper and Darin Acosta. CMS Technical Design Report for the Level-1 Trigger Upgrade. 2013. cds:CMS-TDR-012.

[111] A. Zabi et al. Triggering on electrons, jets and tau leptons with the CMS upgraded calorimeter trigger for the LHC RUN II. *JINST*, 11(02):C02008, 2016. doi:10.1088/1748-0221/11/02/C02008.

[112] P. Sphicas. CMS: The TriDAS project. Technical design report, Vol. 2: Data acquisition and high-level trigger. 2002. cds:CERN-LHCC-2002-026.

[113] CMS Collaboration. CMS Luminosity Measurements for the 2016 Data Taking Period. 2017. cds:CMS-PAS-LUM-17-001.

[114] S. van der Meer. Calibration of the Effective Beam Height in the ISR. 1968. cds:CERN-ISR-PO-68-31.

[115] Martin Breidenbach, Jerome I. Friedman, Henry W. Kendall, Elliott D. Bloom, D. H. Coward, H. C. DeStaebler, J. Drees, Luke W. Mo, and Richard E. Taylor. Observed Behavior of Highly Inelastic electron-Proton Scattering. *Phys. Rev. Lett.*, 23:935–939, 1969. slac-pub-0650.

[116] J. D. Bjorken. Asymptotic Sum Rules at Infinite Momentum. *Phys. Rev.*, 179:1547–1553, 1969. slac-pub-0510.

[117] Andy Buckley et al. LHAPDF6: parton density access in the LHC precision era. *Eur. Phys. J.*, C75:132, 2015. arXiv:1412.7420.

[118] Yuri L. Dokshitzer. Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics. *Sov. Phys. JETP*, 46:641–653, 1977. [Zh. Eksp. Teor. Fiz.73,1216(1977)].

[119] V. N. Gribov and L. N. Lipatov. Deep inelastic e p scattering in perturbation theory. *Sov. J. Nucl. Phys.*, 15:438–450, 1972. [Yad. Fiz.15,781(1972)].

[120] Guido Altarelli and G. Parisi. Asymptotic Freedom in Parton Language. *Nucl. Phys.*, B126:298–318, 1977.

[121] F. D. Aaron et al. Combined Measurement and QCD Analysis of the Inclusive e+- p Scattering Cross Sections at HERA. *JHEP*, 01:109, 2010. arXiv:0911.0884.

[122] S. D. Drell and Tung-Mow Yan. Partons and their Applications at High-Energies. *Annals Phys.*, 66:578, 1971. slac-pub-0808.

[123] John C. Collins. Choosing the renormalization / factorization scale. *J. Phys.*, G17:1547–1549, 1991. doi:10.1088/0954-3899/17/10/009.

[124] Nicholas Metropolis and S. Ulam. The monte carlo method. *Journal of the American Statistical Association*, 44(247):335–341, 1949. doi:10.1080/01621459.1949.10483310.

[125] Richard D. Ball et al. Parton distributions for the LHC Run II. *JHEP*, 04:040, 2015. arXiv:1410.8849.

[126] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014. arXiv:1405.0301.

[127] Paolo Nason. A new method for combining NLO QCD with shower Monte Carlo algorithms. *JHEP*, 11:040, 2004. arXiv:hep-ph/0409146.

[128] Stefano Frixione, Paolo Nason, and Carlo Oleari. Matching NLO QCD computations with parton shower simulations: the POWHEG method. *JHEP*, 11:070, 2007. arXiv:0709.2092.

[129] Simone Alioli, Paolo Nason, Carlo Oleari, and Emanuele Re. A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX. *JHEP*, 06:043, 2010. arXiv:1002.2581.

[130] Johan Alwall et al. Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions. *Eur. Phys. J.*, C53:473–500, 2008. arXiv:0706.2569.

[131] Rikkert Frederix and Stefano Frixione. Merging meets matching in MC@NLO. *JHEP*, 12:061, 2012. arXiv:1209.6215.

[132] Keith Hamilton, Paolo Nason, and Giulia Zanderighi. MINLO: Multi-Scale Improved NLO. *JHEP*, 10:155, 2012. arXiv:1206.3572.

[133] Tom Melia, Paolo Nason, Raoul Rontsch, and Giulia Zanderighi. W+W-, WZ and ZZ production in the POWHEG BOX. *JHEP*, 11:078, 2011. arXiv:1107.5051.

[134] Paolo Nason and Giulia Zanderighi. W^+W^- , WZ and ZZ production in the POWHEG-BOX-V2. *Eur. Phys. J.*, C74(1):2702, 2014. arXiv:1311.1365.

[135] Federico Granata, Jonas M. Lindert, Carlo Oleari, and Stefano Pozzorini. NLO QCD+EW predictions for HV and HV +jet production including parton-shower effects. *JHEP*, 09:012, 2017. arXiv:1706.03522.

[136] John M. Campbell, R. Keith Ellis, and Ciaran Williams. Vector boson pair production at the LHC. *JHEP*, 07:018, 2011. arXiv:1105.0020.

[137] Olivier Mattelaer and Eleni Vryonidou. Dark matter production through loop-induced processes at the LHC: the s-channel mediator case. *Eur. Phys. J. C*, 75:436, 2015. arXiv:1508.00564.

[138] M. Backović, M. Krämer, F. Maltoni, A. Martini, K. Mawatari, and M. Pellen. Higher-order QCD predictions for dark matter production at the LHC in simplified models with s-channel mediators. *Eur. Phys. J. C*, 75:482, 2015. arXiv:1508.05327.

[139] Matthias Neubert, Jian Wang, and Cen Zhang. Higher-order QCD predictions for dark matter production in mono-Z searches at the LHC. *JHEP*, 02:082, 2016. arXiv:1509.05785.

[140] Mikael Chala, Felix Kahlhoefer, Matthew McCullough, Germano Nardini, and Kai Schmidt-Hoberg. Constraining Dark Sectors with Monojets and Dijets. *JHEP*, 07:089, 2015. arXiv:1503.05916.

[141] Torbjörn Sjöstrand et al. An Introduction to PYTHIA 8.2. *Comput. Phys. Commun.*, 191:159–177, 2015. arXiv:1410.3012.

[142] Stefan Ask. Simulation of Z plus graviton/unparticle production at the LHC. *Eur. Phys. J. C*, 60:509, 2009. arXiv:0809.4750.

[143] Massimiliano Grazzini, Stefan Kallweit, Dirk Rathlev, and Marius Wiesemann. $W^{\pm}Z$ production at hadron colliders in NNLO QCD. *Phys. Lett. B*, 761:179, 2016. arXiv:1604.08576.

[144] Julien Baglio, Le Duc Ninh, and Marcus M. Weber. Massive gauge boson pair production at the LHC: a next-to-leading order story. *Phys. Rev.*, D88:113005, 2013. [Erratum: Phys. Rev.D94,no.9,099902(2016)], arXiv:1307.4331.

[145] Massimiliano Grazzini, Stefan Kallweit, and Dirk Rathlev. ZZ production at the LHC: fiducial cross sections and distributions in NNLO QCD. *Phys. Lett. B*, 750:407, 2015. arXiv:1507.06257.

[146] A. Bierweiler, T. Kasprzik, and J. H. Kahn. Vector-boson pair production at the LHC to $\mathcal{O}(\alpha^3)$ accuracy. *JHEP*, 12:071, 2013. arXiv:1305.5402.

[147] Stefan Gieseke, Tobias Kasprzik, and Johann H. Kuhn. Vector-boson pair production and electroweak corrections in HERWIG++. *Eur. Phys. J. C*, 74:2988, 2014. arXiv:1401.3964.

[148] Stefan Höche. Introduction to parton-shower event generators. In *Proceedings, Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders (TASI 2014): Boulder, Colorado, June 2-27, 2014*, pages 235–295, 2015. arXiv:1411.4085.

[149] John M. Campbell, J. W. Huston, and W. J. Stirling. Hard Interactions of Quarks and Gluons: A Primer for LHC Physics. *Rept. Prog. Phys.*, 70:89, 2007. arXiv:hep-ph/0611148.

[150] M. Bahr et al. Herwig++ Physics and Manual. *Eur. Phys. J.*, C58:639–707, 2008. arXiv:0803.0883.

[151] Stefan Höche and Stefan Prestel. The midpoint between dipole and parton showers. *Eur. Phys. J.*, C75(9):461, 2015. arXiv:1506.05057.

[152] CMS Collaboration. Event generator tunes obtained from underlying event and multiparton scattering measurements. *Eur. Phys. J.*, C76(3):155, 2016. arXiv:1512.00815.

[153] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003. doi:10.1016/S0168-9002(03)01368-8.

[154] John Allison et al. Geant4 developments and applications. *IEEE Trans. Nucl. Sci.*, 53:270, 2006. doi:10.1109/TNS.2006.869826.

[155] J. Allison et al. Recent developments in GEANT4. *Nucl. Instrum. Meth.*, A835:186–225, 2016. doi:10.1016/j.nima.2016.06.125.

[156] Stefan Piperov. Geant4 validation with CMS calorimeters test-beam data. In *Hadron collider physics. Proceedings, 19th Symposium, HCP2008, Galena, USA, May 27-31, 2008*, 2008. FERMILAB-CONF-08-371-CMS.

[157] Sunanda Banerjee. Validation of Geant4 Physics Models with LHC Collision Data. Technical Report CMS-CR-2011-038, CERN, Geneva, Feb 2011. cds:CMS-CR-2011-038.

[158] Sunanda Banerjee and Mike D Hildreth. Validation and Tuning of the CMS Simulation Software. Technical Report CMS-CR-2011-019, CERN, Geneva, Jan 2011. cds:CMS-CR-2011-019.

[159] CMS Collaboration. Particle-flow reconstruction and global event description with the CMS detector. *JINST*, 12(10):P10003, 2017. arXiv:1706.04965.

[160] CMS Collaboration. Description and performance of track and primary-vertex reconstruction with the CMS tracker. *JINST*, 9(10):P10009, 2014. arXiv:1405.6569.

[161] CMS Collaboration. Alignment of the CMS tracker with LHC and cosmic ray data. *JINST*, 9:P06009, 2014. arXiv:1403.2286.

[162] R. Fruhwirth. Application of Kalman filtering to track and vertex fitting. *Nucl. Instrum. Meth.*, A262:444–450, 1987. doi:10.1016/0168-9002(87)90887-4.

[163] R. Mankel. A Concurrent track evolution algorithm for pattern recognition in the HERA-B main tracking system. *Nucl. Instrum. Meth.*, A395:169–184, 1997. doi:10.1016/S0168-9002(97)00705-5.

[164] W. Adam, R. Fruhwirth, A. Strandlie, and T. Todorov. Reconstruction of electrons with the Gaussian sum filter in the CMS tracker at LHC. *eConf*, C0303241:TULT009, 2003. arXiv:physics/0306087.

[165] CMS Collaboration. Performance of Electron Reconstruction and Selection with the CMS Detector in Proton-Proton Collisions at sqrt(s) = 8 TeV. *JINST*, 10(06):P06005, 2015. arXiv:1502.02701.

[166] G. L. Bayatian et al. CMS Physics: Technical Design Report Volume 1: Detector Performance and Software. 2006. cds:CMS-TDR-8-1.

[167] CMS Collaboration. Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV. *Submitted to: JINST*, 2018. arXiv:1804.04528.

[168] CMS Collaboration. The performance of the CMS muon detector in proton-proton collisions at sqrt(s) = 7 TeV at the LHC. *JINST*, 8:P11002, 2013. arXiv:1306.6905.

[169] CMS Collaboration. Performance of CMS muon reconstruction in *pp* collision events at $\sqrt{s} = 7$ TeV. *JINST*, 7:P10002, 2012. arXiv:1206.4071.

[170] K. Rose. Deterministic annealing for clustering, compression, classification, regression, and related optimization problems. *Proceedings of the IEEE*, 86(11):2210–2239, Nov 1998. doi:10.1109/5.726788.

[171] CMS Collaboration. Observation of electroweak production of same-sign W boson pairs in the two jet and two same-sign lepton final state in proton-proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. Lett.*, 120(8):081801, 2018. arXiv:1709.05822.

[172] Muon Identification and Isolation efficiency on full 2016 dataset. Technical report, Mar 2017. cds:CMS-DP-2017-007.

[173] Electron and photon performance in CMS with the full 2016 data sample. Technical report, Mar 2017. cds:CMS-DP-2017-004.

[174] CMS Collaboration. Performance of Photon Reconstruction and Identification with the CMS Detector in Proton-Proton Collisions at sqrt(s) = 8 TeV. *JINST*, 10(08):P08010, 2015. arXiv:1502.02702.

[175] Gavin P. Salam. Towards Jetography. *Eur. Phys. J.*, C67:637–686, 2010. arXiv:0906.1833.

[176] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The Anti-k(t) jet clustering algorithm. *JHEP*, 04:063, 2008. arXiv:0802.1189.

[177] Pileup Removal Algorithms. Technical Report CMS-PAS-JME-14-001, CERN, Geneva, 2014. cds:CMS-PAS-JME-14-001.

[178] Jet algorithms performance in 13 TeV data. Technical Report CMS-PAS-JME-16-003, CERN, Geneva, 2017. cds:CMS-PAS-JME-16-003.

[179] CMS Collaboration. Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV. *JINST*, 12(02):P02014, 2017. arXiv:1607.03663.

[180] Jet energy scale and resolution performance with 13 TeV data collected by CMS in 2016. Technical report, Jun 2018. cds:CMS-DP-2018-028.

[181] CMS Collaboration. Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV. *JINST*, 13(05):P05011, 2018. arXiv:1712.07158.

[182] CMS Collaboration. Reconstruction and identification of τ lepton decays to hadrons and ν_τ at CMS. *JINST*, 11(01):P01019, 2016. arXiv:1510.07488.

[183] Performance of reconstruction and identification of tau leptons in their decays to hadrons and tau neutrino in LHC Run-2. Technical Report CMS-PAS-TAU-16-002, CERN, Geneva, 2016. cds:CMS-PAS-TAU-16-002.

[184] CMS Collaboration. Performance of the CMS missing transverse momentum reconstruction in pp data at $\sqrt{s} = 8$ TeV. *JINST*, 10(02):P02006, 2015. arXiv:1411.0511.

[185] Performance of missing transverse momentum in pp collisions at sqrt(s)=13 TeV using the CMS detector. Technical Report CMS-PAS-JME-17-001, CERN, Geneva, 2018. cds:CMS-PAS-JME-17-001.

[186] J. S. Conway. Incorporating Nuisance Parameters in Likelihoods for Multisource Spectra. In *Proceedings, PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding, CERN,Geneva, Switzerland 17-20 January 2011*, pages 115–120, 2011. arXiv:1103.0354.

[187] F. James and M. Roos. Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations. *Comput. Phys. Commun.*, 10:343–367, 1975. doi:10.1016/0010-4655(75)90039-9.

[188] Wouter Verkerke and David P. Kirkby. The RooFit toolkit for data modeling. *eConf*, C0303241:MOLT007, 2003. arXiv:physics/0306116.

[189] Thomas Junk. Confidence level computation for combining searches with small statistics. *Nucl. Instrum. Meth.*, A434:435–443, 1999. arXiv:hep-ex/9902006.

[190] Alexander L. Read. Presentation of search results: The CL(s) technique. *J. Phys.*, G28:2693–2704, 2002. doi:10.1088/0954-3899/28/10/313.

[191] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J.*, C71:1554, 2011. [Erratum: Eur. Phys. J.C73,2501(2013)] arXiv:1007.1727.

[192] S. Dittmaier et al. Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables. 2011. arXiv:1101.0593.

[193] Jon Butterworth et al. PDF4LHC recommendations for LHC Run II. *J. Phys. G*, 43:023001, 2016. arXiv:1510.03865.

[194] B. Delaunay. Sur la sphère vide. A la mémoire de Georges Voronoï. *Bulletin de l'Académie des Sciences de l'URSS. Classe des sciences mathématiques et naturelles.*, 6:793, 1934. MathNet:izv4937.

[195] G. Angloher et al. Results on light dark matter particles with a low-threshold CRESST-II detector. *Eur. Phys. J. C*, 76:25, 2016. arXiv:1509.01515.

[196] R. Agnese et al. New results from the search for low-mass weakly interacting massive particles with the CDMS low ionization threshold experiment. *Phys. Rev. Lett.*, 116:071301, 2016. arXiv:1509.02448.

[197] Xiangyi Cui et al. Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment. *Phys. Rev. Lett.*, 119(18):181302, 2017. arXiv:1708.06917.

[198] D. S. Akerib et al. Results from a search for dark matter in the complete LUX exposure. *Phys. Rev. Lett.*, 118(2):021303, 2017. arXiv:1608.07648.

[199] E. Aprile et al. First Dark Matter Search Results from the XENON1T Experiment. *Phys. Rev. Lett.*, 119(18):181301, 2017. arXiv:1705.06655.

[200] E. Behnke et al. Final results of the PICASSO dark matter search experiment. *Astropart. Phys.*, 90:85, 2017. arXiv:1611.01499.

[201] Search for invisible decays of the Higgs boson produced through vector boson fusion at $\sqrt{s} = 13$ TeV. Technical Report CMS-PAS-HIG-17-023, CERN, Geneva, 2018. cds:CMS-PAS-HIG-17-023.

[202] A. M. Sirunyan et al. Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13\,\text{TeV}$. *Phys. Rev. D*, 97(9):092005, 2018. arXiv:1712.02345.

[203] The CMS Collaboration. Simplified likelihood for the re-interpretation of public CMS results. Technical Report CMS-NOTE-2017-001. CERN-CMS-NOTE-2017-001, CERN, Geneva, Jan 2017. cds:CMS-NOTE-2017-001.

[204] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi. DELPHES 3, A modular framework for fast simulation of a generic collider experiment. *JHEP*, 02:057, 2014. arXiv:1307.6346.

[205] ATLAS Collaboration. Search for an invisibly decaying Higgs boson or dark matter candidates produced in association with a Z boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Lett. B*, 776:318–337, 2018. arXiv:1708.09624.

[206] Chris Quigg. LHC Physics Potential vs. Energy: Considerations for the 2011 Run. 2011. arXiv:1101.3201.