



Search for light scalar particles from Higgs boson decays in exclusive final states with two muons and two hadrons at the CMS experiment

Thesis submitted by Soumya DANSANA

in fulfilment of the requirements of the PhD Degree in Sciences
(ULB - “Doctorat en Sciences”, VUB - “Doctor in de
Wetenschappen”)

Academic year 2025-2026

Supervisors: Professor Barbara CLERBAUX

(Université libre de Bruxelles)

and Professor Steven LOWETTE

(Vrije Universiteit Brussel)

Thesis jury:

Prof. Pascal VANLAER (Université libre de Bruxelles, Chair)

Prof. Alberto MARIOTTI (Vrije Universiteit Brussel, Secretary)

Prof. Markus KLUTE (Karlsruhe Institute of Technology)

Prof. Anna BENECKE (Université catholique de Louvain)

Abstract

The Standard Model of particle physics has been remarkably successful in describing matter and its interactions at the subatomic scale. There are, however, few key observations that do not fit within the framework of the Standard Model, such as the existence and the nature of dark matter. Theoretical extensions, termed as Beyond Standard Model (BSM) theories, often propose new particles that can be produced in high-energy experiments such as CMS at the LHC collider. In this thesis, a search for new scalar particles of $\mathcal{O}(\text{GeV})$ mass in exclusive final states with muons and light hadrons is presented. The search targets exotic decays of the Higgs boson to a pair of prompt or long-lived identical scalar bosons with proper decay lengths up to 100 mm and masses within the range of 0.4–2.0 GeV. This mass window corresponds to a unique phase space where hadronic decays mostly consist of only pairs of light hadrons. The analysis uses proton-proton collision data produced at the LHC in 2016–2018 at a center-of-mass energy of 13 TeV and collected by the CMS detector. The analysis achieves unique sensitivity to very light scalar boson masses and demonstrates a novel approach to probe hadronic decays of light scalar bosons. In parallel, a study of the jet energy resolution at the CMS experiment is carried out using the 2018 dataset to contribute to ongoing jet calibration efforts. A validation study comparing two different methods used within the CMS Collaboration is also performed as part of this work.

Résumé

Le modèle standard de la physique des particules a remarquablement bien réussi à décrire la matière et ses interactions à l'échelle subatomique. Il existe toutefois certaines observations importantes qui ne s'inscrivent pas dans le cadre du modèle standard, comme l'existence et la nature de la matière noire. Les extensions théoriques, appelées théories au-delà du modèle standard (BSM), proposent souvent de nouvelles particules qui peuvent être produites dans des expériences à haute énergie telles que CMS au collisionneur LHC. Dans cette thèse, nous présentons une recherche de nouvelles particules scalaires de masse $\mathcal{O}(\text{GeV})$ dans des états finaux exclusifs avec des muons et des hadrons légers. La recherche cible les désintégrations exotiques du boson de Higgs en une paire de bosons scalaires identiques prompts ou à longue durée de vie, avec des longueurs de désintégration propres pouvant atteindre 100 mm et des masses comprises entre 0,4 et 2,0 GeV. Cette fenêtre de masse correspond à un espace de phase unique où les désintégrations hadroniques consistent principalement en paires d'hadrons légers. L'analyse utilise les données de collisions proton-proton produites au LHC en 2016-2018 à une énergie dans le centre de masse de 13 TeV et collectées par le détecteur CMS. L'analyse atteint une sensibilité unique aux masses de bosons scalaires très légers et démontre une nouvelle approche pour étudier les désintégrations hadroniques des bosons scalaires légers. En parallèle, une étude de la résolution en énergie des jets dans l'expérience CMS est menée à partir de l'ensemble de données de 2018 afin de contribuer aux efforts en cours d'étalonnage des jets. Une étude de validation comparant deux méthodes différentes utilisées au sein de la collaboration CMS est également réalisée dans le cadre de ces travaux.

Samenvatting

Het Standaardmodel van de deeltjesfysica is opmerkelijk succesvol in het beschrijven van materie en haar interacties op subatomaire schaal. Er zijn echter enkele belangrijke waarnemingen die niet binnen het kader van het Standaardmodel passen, zoals het bestaan en de aard van donkere materie. Theoretische uitbreidingen, ook wel Beyond Standard Model (BSM) theorieën genoemd, stellen vaak nieuwe deeltjes voor die kunnen worden geproduceerd in hoge-energie experimenten zoals CMS bij de LHC deeltjesversneller. In dit proefschrift wordt een zoektocht naar nieuwe scalaire deeltjes met een massa van $O(\text{GeV})$ in exclusieve eindtoestanden met muonen en lichte hadronen gepresenteerd. De zoektocht richt zich op exotische vervalprocessen van het Higgs-boson naar een paar identieke scalaire bosonen met een vervalafstand tot 100 mm en massa's in het bereik van 0.4-2.0 GeV. Dit massavenster komt overeen met een unieke faseruimte waar hadronische vervalprocessen voornamelijk bestaan uit paren van lichte hadronen. De analyse maakt gebruik van proton-proton botsingsgegevens die in 2016-2018 bij de LHC zijn geproduceerd bij een massamiddelpuntsenergie van 13 TeV en verzameld door de CMS detector. De analyse bereikt een unieke gevoeligheid voor zeer lichte scalaire bosonmassa's en demonstreert een nieuwe benadering om hadronische vervalprocessen van lichte scalaire bosonen te onderzoeken. Tegelijkertijd wordt een studie uitgevoerd naar de energie-resolutie van jets bij het CMS experiment, waarbij gebruik wordt gemaakt van de dataset van 2018 om bij te dragen aan de lopende studies op het gebied van jetkalibratie. Als onderdeel van dit werk wordt ook een validatiestudie uitgevoerd waarin twee verschillende methoden worden vergeleken die binnen het CMS experiment worden gebruikt.

Acknowledgement

Much like all of written text, the thesis cover carries the name of its author and yet, it would be impossible to not mention the people who have contributed since the beginning and continue to support as this project comes to an end.

I would first like to thank my promoters, Barbara Clerbaux and Steven Lowette, for their guidance, support, and encouragement throughout this PhD. During one of the first discussions, they mentioned that they had previously supervised a PhD together as a team. What I did not expect was just how well they would work together, and I am extremely grateful to have had the opportunity to learn from and interact with both of them over the course of these years. While they complemented each other at every step of this journey, each has shaped my outlook on physics – and beyond – in their own unique way.

To Steven, I would like to say thank you for being an exceptional mentor. Your enthusiasm for physics (and often for debugging!) and the depth to which you are willing to dive into a problem are truly inspiring. I hope to carry forward the spark and excitement for physics that you embody so well. To Barbara, I am grateful for your guidance in navigating the complex endeavour that is both physics and the PhD itself. Your ability to identify what really matters, and to look beyond the thesis has been extremely helpful. There are far too many things I have learned from both of you to list here, but it has been a deeply enriching experience to be guided by you.

I am also grateful to have had the opportunity to work with Laurent Thomas. While this project started as an experimental responsibility towards the CMS ex-

periment, I am glad that while jets did not directly have a role to play in my analysis, I learned so much about it through this project. Thank you for dedicating countless hours to in-depth and fun discussion sessions!

One of the great things about doing a PhD at the IIHE is that there are no bounds on who you can discuss physics problems with. I would like to thank, in particular, Alexander, Jas, Denise, Gerrit, David, AR, Santiago, and Andrea for their valuable insights and support. Within the CMS collaboration, there are too many collaborators to mention individually, but I would like to give special thanks to Lisa, Christina, and Benedikt, who were pivotal in shaping the analysis.

There have been challenges beyond physics which would not be possible to solve without the people at IIHE. This list includes the secretariat at VUB and ULB (in particular, Nina, Sofie, Audrey) for their help with the administrative part and the IT team with helping with maintaining the computing resources.

While this thesis is strictly about physics, the warm and friendly environment of the IIHE has provided me with memories and bonds which I would like to acknowledge. Some of the colleagues who I have shared several fun moments with, include Alexandre, Gerrit, Denise, Marta, Kunal, Eduardo, Felix Heyen, Nordin, Chirayu, Stef, Tiepolo, Kai, Juhee, Jas, Nikos, Godwin, Rijeesh, Marta, Hugues, Max, Elliot, Ilia, Paramita, Jaydeep, Katka, Ali, Felix Schluter, Franco, Golnaz, Arjun, Santiago. I am grateful to have them as colleagues and to have turned many of them into friends. Speaking of friends, I am grateful to have had AR, Indrani and Saranya as flatmates at different points during the last five years. Thank you for making it feel like a home away from home.

I am also deeply grateful to my friends outside the particle physics circle, who have remained close since my undergraduate years. Thank you for providing moments of relief from the stresses and challenges of life. Thank you, Simli, for always being there through both the tough times and the joyful ones, and for keeping me sane throughout this wild journey of the PhD.

Finally, I would like to thank the people who have always been a pillar of support – my parents and my brother. Thank you for your constant belief, encouragement and unwavering support throughout the years.

Soumya Dansana

Contents

Abstract	iii
Introduction	3
1 The SM and extending the Higgs sector	7
1.1 The Standard Model	8
1.1.1 Fermions	8
1.1.2 Quantum Electrodynamics	10
1.1.3 Quantum Chromodynamics	11
1.1.4 Weak interactions and EW theory	12
1.1.5 Brout-Englert-Higgs mechanism	14
1.2 Beyond Standard Model	16
1.2.1 Open questions	16
1.2.2 Extension with light BSM scalars	19
1.2.3 Current collider searches for light particles	26
2 The CMS experiment at the LHC	31
2.1 The Large Hadron Collider	32
2.2 Compact Muon Solenoid experiment	35
2.2.1 CMS coordinate system	37
2.2.2 Silicon tracker	38
2.2.3 Electromagnetic calorimeter	40

2.2.4	Hadronic calorimeter	41
2.2.5	Solenoid Magnet	42
2.2.6	Muon system	43
2.2.7	Trigger system	44
3	Simulation and reconstruction of physics objects	47
3.1	Event simulation	48
3.2	Physics object reconstruction	51
3.2.1	Tracking	51
3.2.2	Primary vertex and vertex reconstruction	54
3.2.3	Particle flow algorithm	55
3.2.4	Electrons and Photons	57
3.2.5	Muons	58
3.2.6	PF Isolation	60
3.2.7	Jets	61
3.2.8	Missing transverse momentum	69
4	Jet energy resolution in the CMS experiment	71
4.1	In-situ method for jet response: p_T -balance	74
4.2	Data and simulated samples	76
4.3	Event selection	77
4.3.1	JER extraction	81
4.3.2	Scale factor derivation	85
4.4	Systematic uncertainties	88
4.5	p_T dependent JER results	92
4.6	Complementary methods: MPF method	95
4.7	Validation of JER scale factors	99
4.7.1	Jet smearing	99
4.7.2	Datasets and Event selection	101
4.7.3	Results	102
4.8	Chapter Summary	104

5 Exclusive search for light scalar bosons	105
5.1 Introduction	105
5.2 Signal simulation and datasets	107
5.2.1 Signal simulation	107
5.2.2 Data and background simulation	114
5.3 Trigger Selection	118
5.4 Event selection	121
5.4.1 Object Selection	121
5.4.2 Analysis Strategy	124
5.4.3 Baseline selection	131
5.4.4 Signal Region & Control Region	132
5.4.5 Reconstructed scalar mass consistency	135
5.4.6 Transverse displacement and Event categories	140
5.4.7 Relative Isolation for muons & hadrons	145
5.4.8 Event selection summary	149
5.5 Background estimation	154
5.6 Full Run 2 data and unblinding	159
5.7 Systematic Corrections and Uncertainties	167
5.7.1 Background Estimation Uncertainties	167
5.7.2 Signal Uncertainties	170
5.8 Results	181
5.8.1 Limit setting procedure	183
5.8.2 Expected and Observed Limits	187
Conclusion and Outlook	195

Introduction

‘What came first, the chicken or the egg?’

This familiar question finds an interesting parallel in particle physics, where theory and experiment have alternately led the way in understanding the fundamental nature of matter and its interactions.

According to a popular particle physics book [1], the inception of particle physics began with discovery of the electron in 1897 by J. J. Thomson [2]. This established the electron as a particle within the atom, taking away the ‘fundamental’ tag from the atom. In this case, experiment preceded the theory. Moving ahead to 1927 when the negative energy states were proposed by Dirac [3], the experiment that confirmed their existence came later in 1932 [4]. Moving ahead to the period between 1964 and 1974, the Standard Model of particle physics was still being constructed. It was proposed by Gell-Mann, Zweig [5, 6] that there were three types of quarks: up, down and strange. However, it was not universally accepted by the community. A fourth quark was proposed by Glashow, Bjorken [7], but was not included in the established picture at the time [1]. Rightly named, the ‘November revolution’ of 1974, marked by the discovery of the J/ψ meson [8, 9], confirmed the existence of a fourth quark. As a result, while a pre-existing theory was proposed beforehand, it was a particle discovery which led to its consolidation.

In the years that followed, the remaining components of the Standard Model were confirmed experimentally. Most of these discoveries came from predictions of the Standard Model, with the last particle of the Standard Model being

confirmed in 2012 [10, 11], the Brout-Englert-Higgs boson or commonly called as Higgs boson. As more open questions continue to emerge that fall outside the scope of the Standard Model, we may again be at the phase where experiment takes the lead and paves the way for the next major breakthrough in particle physics, potentially pointing towards physics Beyond the Standard Model (BSM).

In spite of its great success, the Standard Model has some shortcomings which remain to be addressed, such as the nature of dark matter, the origin of neutrino masses, matter-antimatter asymmetry in the Universe and the large difference in energy scale between gravity and other fundamental forces. A variety of theoretical models have been proposed to provide explanations for these open questions, offering directions to experiments on where to look for new physics.

BSM theories often extend the Standard Model with the addition of new particles. In this thesis, an experimental search is performed with the data collected by the Compact Muon Solenoid (CMS) detector at the Large Hadron Collider (LHC) with one such model that predicts that a new particle can be produced in rare decays of the Higgs boson. In particular, the focus of the analysis is on exploring low-mass or light scalar bosons with masses below 2 GeV by detecting their final-states to two muons or two charged hadrons.

The case described above represents a rare scenario in which the hadronic decay channel of a particle is limited to only two charged hadrons, due to the low-mass of the BSM particle. Above the threshold of 1-2 GeV, the quarks radiate multiple quarks and gluons, leading to the formation of collimated sprays of particles called jets. Since jets are among the commonly used physics objects at the CMS experiment, it is vital to ensure that jets are well-calibrated. Among the various aspects of jet calibration, this thesis focuses on the measurement of the jet energy resolution.

The thesis is structured as follows: Chapter 1 presents a brief overview of the Standard Model, along with the BSM theory that serves as the basis for the BSM physics search and a summary of the current experimental status in this context. Chapter 2 describes the LHC experiment and the CMS detector, outlining the detector's main components and functions. This is followed by Chapter 3 which describes the algorithms and techniques used to reconstruct the collision

event and also provides a brief description of how the collision events are simulated. Chapter 4 presents the methodology and results of the jet energy resolution measurements. Finally, Chapter 5 describes the search for a pair of light scalar bosons produced in Higgs boson decays, where one scalar boson decays to two muons and the other decays to two charged hadrons.

Chapter 1

The Standard Model and extending the Higgs sector

This chapter presents a brief overview of the Standard Model of particle physics, outlining its key principles. It also discusses the theoretical and experimental motivation for the Beyond Standard Model physics search performed in the context of this thesis. The results shown here are taken from various references.

All branches of physics are driven by a close interplay between theory and experiment, and particle physics is no exception. The present theory of fundamental particles and their interactions, the Standard Model (SM), was developed during the 1970s as a result of a series of particle discoveries that began in the 1930s. The model also predicted the existence of several new particles that were eventually discovered. The final missing piece, the Brout-Englert-Higgs boson (BEH boson) or commonly known as the Higgs boson, was discovered in 2012 by the CMS and ATLAS experiments [10, 11], nearly 50 years after its foundations were laid in the Standard Model [12–14].

While it continues to be the guiding framework for particle physics, there are some natural phenomena which still lie beyond its scope. With the aim of addressing these open questions, several Beyond Standard Model (BSM) theories have been proposed, many of which continue to provide the basis of a wide variety of experimental searches for new particles.

In this regard, the text presented in this chapter provides a brief overview of the theoretical aspects relevant to particle physics and, in particular, the experimental work carried out in this thesis. Section 1.1 touches upon the Standard Model and a few of its unsolved challenges, drawing heavily from standard particle physics textbooks [1, 15]. This is followed by discussion of a BSM theory in Section 1.2, which expands upon the results in Refs. [16–18]. A brief highlight of the current experimental results relevant to this BSM theory is provided as well.

1.1 The Standard Model

The Standard Model of particle physics consists of matter particles, termed as fermions, and describes the fundamental interactions between them. Three of the four fundamental forces in nature, namely, the electromagnetic force, the weak force and the strong force, are described by the Standard Model and have associated particle mediators that are termed as gauge bosons. In addition to these, there is the Higgs boson which emerges as a result of the mechanism that allows particles to acquire mass within the framework of the Standard Model. A compact representation of these particles is shown in Fig. 1.1 along with some properties such as their mass and electric charge. They will be discussed more in detail in the following sections.

The mathematical framework of the Standard Model is called a Quantum Field Theory (QFT). In this framework, particles appear as excitations of quantized fields. The fields here are the relativistic extensions of the wavefunction $\Psi(x)$ where x can be any point in space and time (defined as $x = (x_t, x_1, x_2, x_3)$). Similar to quantum mechanics, the dynamics and the interactions of the fields are described by the Lagrangian density, \mathcal{L} .

1.1.1 Fermions

The fermions are spin- $\frac{1}{2}$ particles that obey Fermi-Dirac statistics. In QFT, they are represented by spinor fields, ψ . In the absence of interactions (i.e. free state),

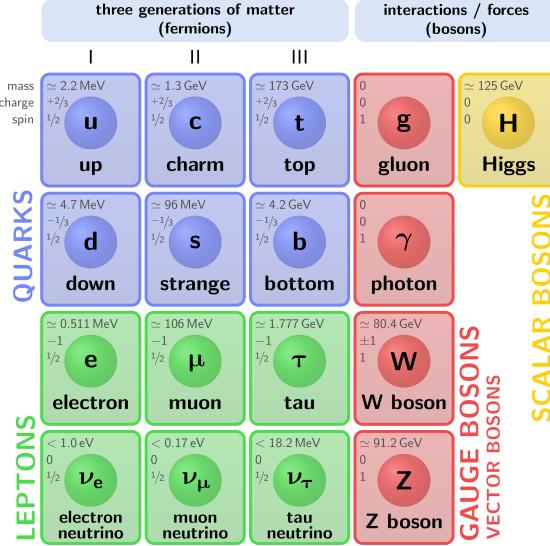


Figure 1.1: A graphical overview of the particles in the Standard Model. Figure taken from Ref. [19].

they are governed by the Dirac Lagrangian which is shown in Eqn. 1.1.

$$\mathcal{L} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi \quad (1.1)$$

Here, μ denotes the spacetime index (running from 0 to 3) and correspondingly, the γ^μ are the Dirac matrices. $\bar{\psi}$ is the adjoint spinor that is defined as $\psi^\dagger \gamma^0$. The term m is the physical mass of the fermion. The Euler-Lagrange (EL) equations for the Dirac Lagrangian eventually yield the equations of motion for a free fermion. The spinor ψ has four components, and the solutions of the resulting equations of motion give two positive energy solutions and two negative energy solutions for the spin- $\frac{1}{2}$ fermion. This was an important theoretical prediction which suggested the presence of antiparticles, SM fermion partners with the same mass and spin but with opposite quantum numbers, such as electric charge.

Particles and antiparticles form the collection of fermions that make up matter in the Standard Model. In nature, there are two broad categories of fermions – quarks and leptons. Both fermions possess intrinsic properties such as electric

charge but the key difference between them is that leptons do not interact via the strong force while the quarks interact dominantly through it.

There are three generations of fermions in nature, as shown in Fig. 1.1. Each generation contains a lepton (e , μ or τ) with charge $-e = -1.602 \cdot 10^{-19}$ C and a corresponding electrically neutral lepton, neutrino (ν_e , ν_μ or ν_τ), which interacts via the weak force only. Additionally, each generation consists of an up and down type quark with electric charge $+2/3 e$ and $-1/3 e$. The first generation quarks are the up (u) and down (d) quarks followed by strange (s), charm (c), bottom (b) and top (t) quarks in the second and third generation respectively.

Across the generations, the mass of the particles increases, as shown in Fig. 1.1. Since only heavier particles can decay to lighter states, the first generation particles are the most stable particles. Considering the simplest case of no off-shell mediator, or phase space suppression, the lifetime can be approximated as $\tau \sim 1/m$ [20]. Accordingly, the lifetime τ of the particles and consequently the stability decreases for the second and third generations. As a reference, the heaviest SM fermion, the top quark, has an extremely short lifetime of $\sim 10^{-25}$ seconds.

1.1.2 Quantum Electrodynamics

In QFT, the interactions emerge as a result of underlying symmetries and their respective transformation properties. For instance, the Lagrangian in Eqn. 1.1 is constructed in the context of special relativity where the equations of motions remain invariant under Lorentz transformations. This translational invariance is a global symmetry which, according to Noether's theorem, is associated to a conserved quantity in the system. In this case, the translational invariance leads to the conservation of momentum.

Consider a global phase rotation characterized by the transformation, $\psi(x) \rightarrow \psi' = e^{i\alpha} \psi(x)$. This is called a U(1) transformation. Under this transformation, the Lagrangian in Eqn. 1.1 remains invariant and has an associated U(1) symmetry. However, if the transformation is a local phase rotation (i.e. a function of space-time coordinate), then the transformation $\psi(x) \rightarrow \psi' = e^{i\alpha(x)} \psi(x)$ does not maintain the invariance of the Lagrangian.

To ensure that the gauge invariance is maintained, the Lagrangian is modi-

fied according to Eqn. 1.2 with the introduction of a new vector field, A_μ , that transforms as Eqn. 1.4 to cancel any phase terms (i.e. $\alpha(x)$ dependent terms).

$$\mathcal{L} = \bar{\psi}(i\gamma^\mu(\partial_\mu - igA_\mu) - m)\psi \quad (1.2)$$

$$\equiv \bar{\psi}(i\gamma^\mu \mathcal{D}_\mu - m)\psi \quad (1.3)$$

$$A_\mu \rightarrow A'_\mu = A_\mu + \frac{1}{g}[\partial_\mu \alpha(x)] \quad (1.4)$$

Here, the partial derivative ∂_μ is replaced with the covariant derivative, \mathcal{D}_μ , while g is the interaction strength between the field ψ and gauge field A_μ . With the inclusion of gauge invariant kinetic terms for A_μ , the new Lagrangian is shown in Eqn. 1.5 where $F_{\mu\nu}$ is the field strength tensor, defined as $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$.

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \bar{\psi}(i\gamma^\mu \mathcal{D}_\mu - m)\psi \quad (1.5)$$

Equation. 1.5 therefore shows the Lagrangian for a generic U(1) gauge theory. The electromagnetic interaction follows a $U(1)_{EM}$ gauge theory and is termed as quantum electrodynamics (QED). The particle corresponding to the gauge field, A_μ , is the gauge boson, the photon, and is the mediator particle. Note that the addition of a mass term for the photon would break gauge invariance and hence, the photon remains massless in this theory.

In a gauge theory, for each degree of freedom, there is an associated operator which is responsible for the transformation. These operators are called the generators of the gauge theory. QED has only one degree of freedom and the associated generator is the electric charge operator, Q . This relates to the associated quantum number in the theory when acting on a particle, which in this case is the electric charge.

1.1.3 Quantum Chromodynamics

With a similar approach as shown before, gauge theories with larger number of degrees of freedom are developed by introducing new gauge fields to describe the other interactions of the Standard Model. Quantum Chromodynamics (QCD) is a $SU(3)_C$ gauge theory that describes the strong nuclear force. The theory has

8 degrees of freedom and therefore has 8 generators, typically represented by 3×3 Gell-Mann matrices λ^a ($a = 1, \dots, 8$). Correspondingly, there are 8 massless gauge boson fields G_μ^a called gluons. These fields also modify the definition of the covariant derivative, \mathcal{D}_μ accordingly.

Quarks carry with them the colour charge, which appears as a three-element vector. The associated colour charges are r, b, g and their anticolours. The gluons, which mediate the interaction between the colour fields \bar{q} and q (Eqn. 1.1 with 3-component vectors for the Dirac fields), have 8 possible linear colour-anticolour combinations. In nature, only particles with zero net colour charge (referred to as colour singlets) have been observed freely. This leads to the concept of colour confinement, which means that quarks and gluons always appear in colour-neutral states called hadrons.

The coupling strength for QCD is given by the strong coupling constant $g_s = \sqrt{4\pi\alpha_s}$. While this is denoted as a constant, it has been observed that α_s is a function of the energy scale or the momentum transfer. As the energy scale of the interaction increases, α_s decreases. Above a certain energy scale, the strong nuclear interactions are well described by perturbative QCD (pQCD). At the LHC, where high-energy pp collision takes place, pQCD is important in describing how the colliding partons (quarks and gluons) interact and radiate more quarks and gluons through parton showers (described in Section 3.1). As they approach the non-perturbative regime, the colour confinement processes dominate and leads to stable hadrons through hadronization.

Since the pQCD methods are not applicable at low energies, effective theories are often employed to describe the dynamics of hadrons. One such example is the Chiral Perturbation theory which is briefly described in Section 1.2.2.

1.1.4 Weak interactions and EW theory

The weak interactions can be described by the $SU(2)$ gauge group theory. The gauge bosons in this theory allow flavour changing interactions between fermions through charged currents to take place. An example of this is the decay of a muon to an electron. Through experimental observations, it was found that weak interactions through flavour-changing currents exhibit maximal parity violation, indicating that they are chiral in nature. Within this structure, only the left-handed

fermions form $SU(2)$ doublets while the right-handed fermions appear as singlets and do not participate in flavour-changing weak interactions. In addition to this, there are neutral current weak interactions as well, which do not change the flavour of the fermions.

After restricting the $SU(2)$ structure to only left-handed fermions (termed the $SU(2)_L$ gauge theory), there was still the need to find the appropriate symmetry which is conserved in this structure. The $SU(2)_L$ gauge group consists of 3 gauge bosons, $W_\mu^1, W_\mu^2, W_\mu^3$ with the corresponding charge being the weak isospin (T_3). However, this weak isospin charge was not found to be conserved in the interactions while the electric charge conservation was still followed. As a result, a unified theory, the $SU(2)_L \times U(1)_Y$ [21–23] was developed.

In this unified description, the electroweak theory, the $U(1)_Y$ gauge group introduces one gauge field, B_μ and the corresponding quantum number, the weak hypercharge, Y . Since $U(1)_Y$ describes the same electromagnetic physics as $U(1)_{EM}$, it follows the same mathematical structure. However, $U(1)_{EM}$ is the residual symmetry after the $SU(2)_L \times U(1)_Y$ symmetry is broken through spontaneous symmetry breaking, which is presented later. The conserved quantity, the electric charge, is calculated as $Q = T_3 + Y/2$.

The physical observable fields of the electroweak theory arise from the superposition of the four gauge bosons listed above. These are listed in Eqn. 1.6, 1.7, 1.8 and correspond to the two charged W boson fields W_μ^\pm , the neutral Z boson field Z_μ^0 and the photon field A_μ . θ_W denotes the weak mixing angle.

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \quad (1.6)$$

$$Z_\mu^0 = \cos \theta_W W_\mu^3 - \sin \theta_W B_\mu \quad (1.7)$$

$$A_\mu = \sin \theta_W W_\mu^3 + \cos \theta_W B_\mu \quad (1.8)$$

Note that in this theory as well, all the gauge bosons are massless. However, it is observed in experiments that the W^\pm and Z bosons are massive particles.

1.1.5 Brout-Englert-Higgs mechanism

The final piece of the puzzle is the BEH mechanism which allows explicit mass terms of the gauge bosons and the fermions to enter the SM Lagrangian without violating gauge invariance. The formalism developed by Robert Brout, François Englert and Peter Higgs [12–14] introduces a complex scalar field ϕ to the $SU(3)_C \times SU(2)_L \times U(1)_Y$ theory developed so far. Since the QCD part is unaffected by this, we focus on the electroweak part. The scalar field ϕ is a doublet under $SU(2)_L$ and has a weak hypercharge of $Y = 1$. It is represented as Eqn. 1.9

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix} \quad (1.9)$$

The scalar follows the Klein-Gordon Lagrangian, shown in Eqn. 1.10 with a potential $V(\phi)$ (Eqn. 1.11). The parameters, μ^2 and λ are real arbitrary parameters that define the shape of the potential.

$$\mathcal{L} = (D^\mu \phi)^\dagger (D_\mu \phi) - V(\phi) \quad (1.10)$$

$$V(\phi) = -\mu^2(\phi^\dagger) + \lambda(\phi^\dagger \phi)^2 \quad (1.11)$$

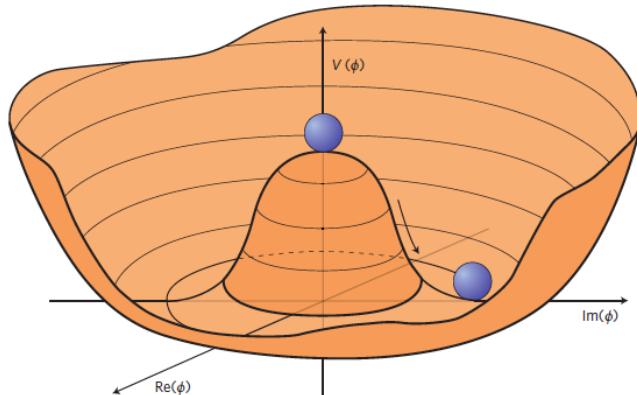


Figure 1.2: Illustration of the Higgs boson potential. Figure taken from Ref. [24].

When $\mu^2, \lambda > 0$, the shape of the potential is shown in Fig. 1.2. The potential has a degenerate set of minima, where the value of the potential is non-zero.

Visually, the set of minima lie on the circle in the $Re(\phi)$ and $Im(\phi)$ plane at the lowest possible $V(\phi)$. Accordingly, the potential is symmetric under rotations in $Re(\phi)$ and $Im(\phi)$ but would change under radial transformations. The minimal value, referred to as the vacuum expectation value (VEV) $v = \sqrt{\frac{\mu^2}{2\lambda}}$, is measured to be ~ 246 GeV.

The next step is gauge fixing which essentially means choosing a particular configuration from the set of minima and breaking the degeneracy with a small perturbation, $h(x)$. This is shown in Eqn. 1.12.

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} \quad (1.12)$$

Performing this gauge fixing eventually leads to the masses of the Higgs boson, the W^\pm boson and the Z boson in terms of v , λ , g_1 , g_2 (coupling constants of $SU(2)_L \times U(1)_Y$) and θ_W .

Finally, the Yukawa Lagrangian (Eqn. 1.13) is introduced which couples a right-handed fermion field with its left-handed counterpart (and its hermitian conjugate or h.c.) through the scalar field ϕ . After taking the VEV for the scalar field, fermion masses arise which are directly proportional to the yukawa couplings, Y_{*j} .

$$\begin{aligned} \mathcal{L}_{Yukawa} = & \sum_{j=1}^3 Y_{d_j} \bar{Q}_{jL} \phi d_{jR} + (h.c.) \quad (\text{d, s, b mass term}) \\ & \sum_{j=1}^3 Y_{u_j} \bar{Q}_{jL} \phi u_{jR} + (h.c.) \quad (\text{u, c, t mass term}) \quad (1.13) \\ & + \sum_{j=1}^3 Y_{l_j} \bar{l}_{jL} \phi l_{jR} + (h.c.) \quad (\text{e, } \mu, \tau \text{ mass term}) \end{aligned}$$

At this stage, the BEH mechanism is complete, giving mass to the gauge bosons and the fermions while preserving gauge invariance under the $SU(3)_C \times SU(2)_L \times U(1)_Y$ structure of the Standard Model. A similar mechanism involving a spontaneous symmetry breaking procedure is described in the next section in the context of a Standard Model extension with an additional new scalar field.

1.2 Beyond Standard Model

1.2.1 Open questions

The Standard Model has been extremely successful in describing the fundamental nature of matter and its interactions through electromagnetic, weak and strong forces. It manages to provide a framework that unites three distinct interactions under one umbrella and continues to provide accurate predictions that are tested and confirmed in particle physics experiments. There are however some phenomena which are not accounted for in the Standard Model. Some of them have been briefly described here.

Gravity

As mentioned previously, the Standard Model does not include the gravitational force in its framework. Since the theory of General Relativity describes gravity as the curvature of space-time, there have been several efforts at forming a quantum field theory for gravity. A hypothetical particle associated to the quantum gravity field is commonly referred to as the graviton. Although this extension may seem conceptually straightforward, attempts to formulate it have resulted in a nonrenormalizable theory [25].

Neutrino masses

From the Yukawa Lagrangian in Eqn. 1.13, it follows that while charged leptons and quarks acquire masses through their couplings to the Higgs field, neutrinos remain massless within the Standard Model framework. This is because no right-handed neutrino has been observed, and accordingly, the SM does not have Yukawa coupling term for neutrinos. As a result, neutrinos were considered to be massless. The discovery of neutrino oscillations by the SNO [26] and Super-Kamiokande experiments [27] provided evidence that neutrinos are not massless. Neutrino oscillations refer to the phenomenon where neutrinos change flavour as they travel. This is only possible if the neutrino propagates as a superposition of mass eigenstates with different masses, demonstrating that neutrinos are indeed massive particles.

Dark matter

Over the past several decades, several astrophysical and cosmological experiments and observations have supported the claim that only 5% of the energy content of the universe consists of ordinary matter and falls within the framework of the Standard Model. The remaining portion is attributed to two largely unknown components: dark energy and dark matter. Recent analyses of cosmological data [28] suggests that about 67% of the universe is filled with dark energy while the remaining 33% contains dark matter and ordinary matter. While dark energy continues to be a largely unknown mystery, there have been several studies to understand the nature of dark matter, which is nearly five times more abundant than ordinary matter.

The existence of dark matter was first proposed when several cosmological observations suggested the presence of more mass than what was visible in galaxies (also called luminous matter). Considering the orbital velocity $v = \sqrt{\frac{GM(r)}{r}}$, where $M(r)$ is the enclosed mass of luminous matter in the galaxy and r is the distance from the center of the galaxy, it is expected that as one moves radially away from the center, the mass increases. Since most of the luminous matter is concentrated near the galactic center, after a certain point, the mass should become nearly constant with r . In such a case, it is expected that the velocity of objects is expected to decrease as one moves further away. However, on plotting the velocity as a function of r , it was seen that the velocity becomes nearly constant with the radial distance, as shown in Fig. 1.3. This observation hinted towards the presence of non-luminous matter, now known to be dark matter.

Another observation that supports the existence of dark matter comes from gravitational lensing. As light travels in space, the gravitational pull of massive objects would cause the light to bend in its path, according to General Relativity. This leads to the appearance of the light emitting source to appear as a distorted image of the distant object, possibly as a ring or halo of light around the massive object. This effect is called gravitational lensing. By carrying out observations of distant bright objects, the changes in the intensity of light can be used to infer the presence of these massive objects. It has been observed in several cases that there is a constant background component which is not accounted for by the mass

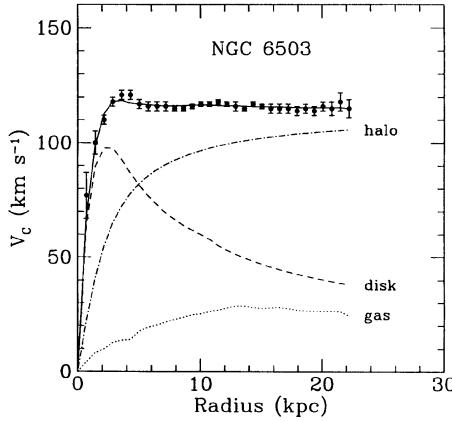


Figure 1.3: Galactic rotation curve for NGC 6503 showing the contributions from the disk (luminous matter) and (insterstellar) gas components. The halo denotes the dark matter contribution needed to reproduce the observed galactic curve. Figures taken from Ref. [29].

of the luminous objects. Other cosmological observations such as the anisotropic structure of Cosmic Microwave Background (CMB), support the existence of dark matter. A review of these observations is presented in Ref. [29, 30].

The hierarchy problem

In a broad sense, the hierarchy problem relates to the large difference in scale between the electroweak theory ($\Lambda_{EW} \sim 10^2$ GeV) and the scale at which quantum gravity effects can show up, the Planck scale ($\Lambda_{Pl} \sim 10^{19}$ GeV). A more technical description of the hierarchy problem arises from the corrections on the Higgs boson mass which is many orders larger than its observed mass of 125 GeV.

In the discussions of the Standard Model presented before, the focus was on studying the Lagrangian and its form. This is followed by calculating the amplitude of an interaction or the matrix element, \mathcal{M} . These calculations are performed perturbatively through Feynman diagrams, providing the results at different order of the couplings based on how many vertices are present in the diagram. The lowest order represents the dominant contribution while larger orders through loop diagrams are included as corrections to the tree-level contribution.

In QFT, divergences often appear in these corrections and through renormalization techniques, a cutoff scale Λ is introduced to quantify the divergence. When the divergent term is logarithmic (of the form $\sim \log \Lambda^2$), the correction factor is small and increases slowly with the energy. An example of such a case is QED, where one-loop corrections lead to mild logarithmic divergences.

However, when considering the correction to the Higgs mass from the one-loop fermion diagram, the divergence is quadratic (of the form $\sim \Lambda^2$). Since the cutoff can be taken as the highest energy scale in the theory, at $\Lambda = \Lambda_{Pl}$, this would amount to a massive correction factor on the mass of the Higgs boson. Therefore, for the Higgs boson mass to be 125 GeV, a large and scale-sensitive cancellation term is required, which appears unnatural. This fine-tuning directly reflects the large hierarchy between the electroweak scale and the Planck scale. Although the observed Higgs boson mass is small and stable, the unnaturally large and precise cancellation required at the Planck scale may hint at the existence of new physics that regulates this fine-tuning.

1.2.2 Extension with light BSM scalars

Since the inception of the Standard Model, there have been hints that there may be some limitations of the Standard Model. These open questions have led to the proposal of many BSM theories. As more pieces of the Standard Model were discovered the theory focus has accordingly evolved. For instance, theories based on supersymmetry [31–37], typically follow a top-down approach, where fundamental symmetries and gauge groups are proposed, forming a complete model with distinct predictions for new particles and interactions. Due to the highly specific nature of these BSM theories, a more general approach has been adopted in recent times through the development of simplified models. These models typically emphasize on developing the phenomenological signatures to guide the experimental searches. They can be interpreted as effective models that describe the interactions of potential BSM particles and are often constructed as minimal extensions of the Standard Model. These particles could arise as mediators between the BSM sector and the SM and couple directly to SM particles.

In the context of collider searches, many simplified models can be placed in four broad categories, based on how they couple with the Standard Model par-

ticles or the ‘portal’ with which they interact with the Standard Model. These consist of the Spin-1 portal, Spin-0 portal, Neutrino portal and the Fermion portal. Each of the portals have their own unique interactions but can also lead to overlapping experimental signatures. Since the discussion on these portals falls outside the scope of the thesis, the reader is instead referred to Ref. [38] for a complete overview.

Instead, the focus in this thesis is a model that falls under the category of the Spin-0 portal. This portal consists of a range of models which propose a scalar or pseudoscalar mediator which couples to BSM theories (such as Dark Matter). A subset of these include scalars which mix with the SM Higgs boson and inherit suppressed couplings to SM fermions. The mass and the eventual coupling of the BSM scalar boson are mostly free parameters, which carry motivations from theory and are constrained through experiments. In recent times, there has been a push from experimental searches towards exploring the parameter space at lower masses and small couplings. In particular, light BSM scalar bosons with mass of $\mathcal{O}(GeV)$ appear in several SM extensions including cosmic inflation models [39–41], supersymmetry [42, 43], dark matter models [44, 45] and cosmological solutions to the hierarchy problem [46].

In this thesis, we focus on the model presented in Ref. [16], a minimal extension of the Standard Model with a single scalar boson of $\mathcal{O}(GeV)$ mass that mixes with the SM Higgs boson. In this section, a brief overview of the model is presented. For an in-depth motivation for this model, the reader is referred to Refs. [16–18]. In the next section, a few experimental searches for BSM particles of $\mathcal{O}(GeV)$ mass are described to illustrate the reach of the current results. Together, these sections provide the motivation for a physics search for light BSM scalar bosons that forms the basis of this thesis and is presented in Chapter 5.

Higgs-scalar mixing

Following the descriptions in Refs. [16, 17], the framework of the BSM theory is outlined here. A minimal extension of the SM is considered by the introduction of one real singlet scalar field S . A general Lagrangian, taking care that the

dimensions are not greater than 4, looks like Eqn. 1.14.

$$\mathcal{L} = \mathcal{L}_{SM} + \alpha_1 S + \alpha_2 S^2 + \alpha_3 S H^\dagger H + \alpha_4 S^2 H^\dagger H + V_{int}(S) \quad (1.14)$$

An approximate \mathbb{Z}_2 symmetry is imposed. The \mathbb{Z}_2 symmetry has the same mathematical structure as the parity symmetry, and performs the following transformation, $S \rightarrow -S$. Since this is a new symmetry, it is assumed that the \mathcal{L}_{SM} is invariant under it. The \mathbb{Z}_2 symmetry forbids terms with odd powers of S in the Lagrangian. The ‘approximate’ nature of it means that any odd powers appear as small perturbations. Using this information, the Eqn. 1.14 can be simplified to Eqn. 1.15.

$$\mathcal{L} = \mathcal{L}_{SM} + \frac{1}{2} \partial_\mu S \partial^\mu S - \frac{1}{2} \tilde{m}_S^2 S^2 + (\mu S + \frac{1}{2} \lambda_{SH} S^2) H^\dagger H - V_{int}(S) \quad (1.15)$$

The parameter, \tilde{m}_S is the mass parameter for the BSM scalar S , while the parameters μ and λ_{SH} are the coupling constants. Following the approximate \mathbb{Z}_2 symmetry, it is expected that the parameter μ is small. This term induces mixing between the SM Higgs boson and the BSM scalar boson, as will be seen after the spontaneous symmetry breaking is performed. Fixing the gauge as listed in Eqn. 1.16, 1.17:

$$\langle H \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} \quad (1.16)$$

$$\langle S \rangle = s_0 + s \quad (1.17)$$

with s_0 as the VEV for the BSM scalar field S and $v = 246$ GeV as the VEV for the SM Higgs field, leads to the calculations similar to the spontaneous symmetry breaking for a standard model Lagrangian. In this scenario, we focus on the portion relevant to describe the interaction of the BSM scalar field.

First, the terms with dimension two (i.e. sh, s^2, h^2) are rewritten as Eqn. 1.18

$$\mathcal{L} \supset \begin{pmatrix} h & s \end{pmatrix} \mathcal{M} \begin{pmatrix} h \\ s \end{pmatrix} \quad (1.18)$$

where \mathcal{M} can be interpreted as the mass squared matrix of the Higgs-singlet system. In the limit of small mixing, this leads to Eqn. 1.19 where m_h^2 is the SM Higgs mass parameter.

$$\mathcal{M} = \begin{pmatrix} m_h^2 & v(\mu + \lambda_{SH}s_0) \\ v(\mu + \lambda_{SH}s_0) & m_S^2 + \frac{1}{2}\lambda_{SH}v^2 + V''_{\text{int}}(s_0) \end{pmatrix} \quad (1.19)$$

The next step is the rotation of the current basis to the mass basis by performing the transformation in Eqn. 1.20

$$\begin{pmatrix} H \\ S \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} h \\ s \end{pmatrix} \quad (1.20)$$

where θ is the mixing parameter between the Higgs and the BSM scalar fields. Under the small mixing assumption, this leads to a simple transformation, as listed in Eqn. 1.21, 1.22.

$$H = h + \theta s \quad (1.21)$$

$$S = s - \theta h \quad (1.22)$$

Plugging the terms after the transformation back into the Lagrangian and focusing on the interaction terms between the BSM scalar boson and SM particles, leads to Eqn. 1.23.

$$\mathcal{L} \supset -\theta \frac{m_f}{v} S f \bar{f} + 2\theta \frac{m_W^2}{v} S W^+ W^- + \theta \frac{m_Z^2}{v} S Z^2 + \lambda_{SH} v \left(\frac{1}{4v} S^2 H^2 + \frac{1}{2} S^2 H \right) \quad (1.23)$$

It is seen that the BSM scalar boson inherits the couplings of the Higgs boson to the SM particles with a suppression factor determined by the mixing parameter, θ (or $\sin(\theta)$ or s_θ in some references). The interaction term, $-\theta \frac{m_f}{v} S f \bar{f}$ describes how the BSM scalar boson couples to fermions. For BSM scalar bosons with masses below the B-meson mass threshold, this interaction term allows the possibility of light BSM scalar bosons with masses below 5 GeV to be produced in rare B-meson decays. On the other hand, the interaction term $\frac{1}{2}\lambda_{SH}vS^2H$,

allows for the decays of the Higgs boson directly to a pair of light BSM scalar bosons. Both of these production modes are useful to search for light BSM scalar bosons and have been pursued in several of the analyses which will be discussed in the next section. For the analysis performed in this thesis, the Higgs-mediated production of the light scalar bosons has been pursued and is well-motivated by the current estimates of the SM Higgs boson properties. This is discussed further in Chapter 5.

So far, the discussion has been centered on building the model which induces a small mixing between a light BSM scalar boson and the SM Higgs boson. While the BSM parameters of the model are free parameters, some constraints on s_θ and λ_{SH} have been calculated in Ref. [17]. These constraints also provide a view of the experimental signatures for these particles. For instance, the BSM light scalar bosons can decay to SM fermions, which can be potentially detected in experiments. However, the lifetime of the BSM scalar boson is expected to be large due to the small s_θ .

A simple toy formula in Ref. [20] and shown in Eqn. 1.24 illustrates how the partial width of the BSM scalar boson depends on the mixing parameter for its decays to SM fermions.

$$\Gamma \sim \frac{\epsilon^2}{8\pi^{a-1}} \frac{m^n}{M^{n-1}} \quad (1.24)$$

Here, ϵ is a small dimensionless parameter such as the mixing angle and m is the mass of the BSM particle. M is the mass of any heavy, off-shell particle that the BSM particle decays through. n is an odd, positive integer which determines the suppression factor related to this type of decay while a is a positive integer, indicating the number of final state particles. For the Higgs-scalar mixing model considered here, the decay width follows the relation in Eqn. 1.25.

$$\Gamma \propto s_\theta^2 m s \quad (1.25)$$

Since $\tau \sim \frac{1}{\Gamma_{\text{tot}}}$, the BSM scalar boson can have a large proper lifetime. Furthermore, as the distance travelled in lab-frames is $L = \beta \gamma c \tau$, it is likely that the SM decays of the scalar bosons can occur at large distances from the production

point. This is especially important because this leads to a displaced signature which is pursued in several physics analyses.

Light scalar boson decays

The main decay modes for the light scalar boson are decays to photons, leptons and hadrons in the $\mathcal{O}(\text{GeV})$ mass range. For the analysis presented in this thesis, the mass range of interest is $m_S \leq 2 \text{ GeV}$. While the decay rate to photons and leptons follows a straightforward calculation, the hadronic channel is a complicated affair due to several energy scales in this mass range that dictate the interaction.

Using Fermi's Golden rule, the rate of a scalar particle decaying to two identical leptons follows from the Lagrangian 1.23. The calculation steps are outlined in Ref. [1] and at tree-level, after accounting for the possible spin-states, leads to Eqn. 1.26

$$\Gamma(S \rightarrow l\bar{l}) = \frac{g^2 m_S}{8\pi} \beta^3 = \frac{s_\theta^2 G_F m_S}{4\sqrt{2}\pi} m_l^2 \beta^3 \quad (1.26)$$

$$\beta = \sqrt{1 - \frac{4m_l}{m_S}} \quad (1.27)$$

where $g = s_\theta \frac{m_f}{v}$ is the coupling from the interaction term and G_F is the Fermi constant ($= \frac{1}{\sqrt{2}v^2}$).

For the hadronic channel, the decay modes and the respective rates are calculated according to the mass of the scalar boson and the cutoff scales, accounting for various regimes of QCD. In general, the decays of the scalar boson to quarks (i.e. $S \rightarrow q\bar{q}$) is described by perturbative QCD when the strong coupling constant, α_S is small. In this regime, the quarks radiate more quarks and gluons, in a process called hadronization (see Section 3.1), leading to a final-state of collimated spray of particles, called jets. However, the non-perturbative effects of QCD start to dominate at low energies and require alternate descriptions. As a reference, the cutoff scale for QCD, below which the non-perturbative regime is entered, is $\Lambda_{\text{QCD}} \sim 1 \text{ GeV}$. In this region, the Chiral Perturbation Theory (or ChPT) calculations are applied to predict the decay width to pions. However, this is not a sharp cutoff, and there exists a transition region between 1–2 GeV where

non-perturbative effects are still present. This region relies on dispersion relations to estimate the hadronic decay rates. At the energy scale of $\Lambda_S^{\text{pert}} \sim 2 - 4$ GeV, perturbative QCD starts to become relevant and describes the decays to quarks and gluons.

Accordingly, the decay rates for the light scalar bosons to hadrons are computed by employing the three methods listed above using various calculations [47–52]. The most recent calculation in Ref. [16] updates/combines them to provide the partial widths which form the basis for the branching ratios calculated in Ref. [17, 18]. Since the detailed calculations are quite involved and complex, only the relevant concepts and results are presented here.

ChPT is an effective field theory for low-energy QCD interactions. This means that the ChPT is valid up to a certain scale which, in this case, is the Λ_{QCD} . For the light quarks, (u,d,s quarks), the small differences between the quark masses allows for them to be treated as approximately massless. In this limit, the theory can then be considered to have two independent global SU(2) chiral symmetries (i.e. $\text{SU}(2)_L \times \text{SU}(2)_R$).

At these low energies, the quarks are confined to hadrons. The quark bilinears, $\bar{q}q$, in this case, can form both scalar and pseudoscalar mesons, as expected from the combinations between the left-handed quarks and right-handed quarks. However, experimentally, the lightest meson octet consists of particles with only odd-parity (i.e. pseudoscalar mesons). This indicates that the $\text{SU}(2)_L \times \text{SU}(2)_R$ symmetry is broken to a single $\text{SU}(2)_{\text{isospin}}$ symmetry. The spontaneous symmetry breaking of the chiral symmetry leads to Goldstone bosons that can be identified as the $\pi^+, \pi^-, \pi^0, \eta, K^+, K^-, K^0, \bar{K}^0$ mesons. By introducing a composite field to represent the mesons and also include the interactions between the quarks, a general (or effective) Lagrangian is built with it. After applying the conditions above, the effective Lagrangian is used to compute matrix elements which eventually are matched to information from known results, like the pion lifetime.

Through the ChPT approach, decay rates for $m_S \leq 1$ GeV are calculated. For $1 < m_S \leq 2$ GeV, dispersion analysis is applied which uses analyticity and unitarity conditions based on complex analysis techniques, to evaluate the decay widths without any prior assumptions on the underlying theory. Using $\pi\pi$ and

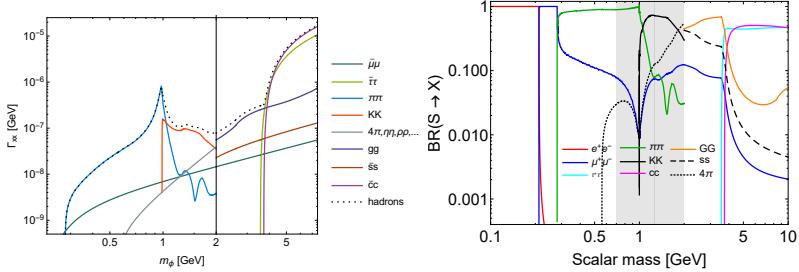


Figure 1.4: The decay width (left) and the branching ratio (right) for decays of light scalar bosons to SM fermions. The decay width for the individual modes is calculated with $s_\theta^2 = 1$. Figures taken from Ref. [16] (left) and Ref. [18] (right).

KK scattering data, the decay widths to pions and kaons are calculated. Some multi-meson channels also have small contributions to the total decay width and become relevant from $m_S \gtrsim 1.2$ GeV. This is derived from a toy model and ensuring continuity of the total decay width at $m_S = 2$ GeV, above which perturbative QCD results are applicable. Since the dispersive method relies heavily on experimental input, there are significant uncertainties associated with it due to the lack of data on meson scattering at high energies.

The decay width and the corresponding branching ratio for the various hadronic and leptonic channels are shown in Fig. 1.4. At $m_S \sim 1$ GeV, the decay channel to kaons is accessible kinematically. This leads to a rise in the branching ratio for the pion and the kaon decay modes but a sharp drop for the branching ratio to other decay modes like muons. The sharp rise can be attributed to the strong overlap and interference effects between the pion and kaon channels.

1.2.3 Current collider searches for light particles

Prior to the discovery of the SM Higgs boson, extensive searches for BSM scalar particles were carried out across a broad range of masses, as the SM Higgs boson mass was unknown. As a result, there have been several probes for light BSM particles that have spanned across multiple experiments (present and past). In this section, some relevant collider searches for light BSM particles are briefly discussed to illustrate the current experimental coverage. Since most of them

aim to be as model agnostic as possible, a range of models have been considered by the analyses. However, in such cases, the final-states that are probed have the potential to also target possible light BSM scalar bosons according to the model discussed above. The emphasis here is on scenarios where a BSM particle with masses below 2 GeV is produced from SM states and decays back into SM states within the detector volume.

Light BSM particles in rare meson decays

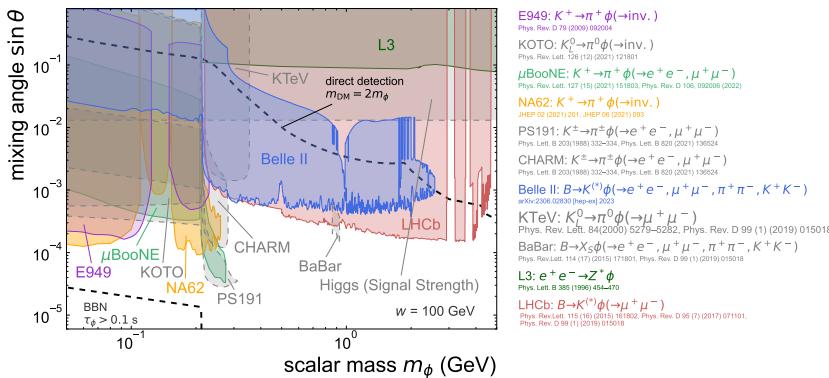


Figure 1.5: An overview of the current upper limits on the mixing angle $\sin \theta$ as a function of the scalar mass, m_ϕ using the results from various experiments. Additionally, the signal strength measurements of the SM Higgs boson provide indirect constraints on the mixing angle as well [53, 54]. Figure is taken from Ref. [54].

Following the Standard Model extension discussed before, light scalar bosons can be produced in rare meson decays. In particular, there are several searches for scalar bosons that have been carried out in decays of kaons and B-mesons, targeting various final-state particles. There have also been searches which target scalar bosons decaying outside the detector volume, using the missing energy or momenta to identify possible signal events.

The sensitivity of the searches are often shown in terms of a two-dimensional map of the free parameter space, (θ, m_ϕ) . A brief summary of the results are discussed in Ref. [54]. Fig. 1.5 shows the scanned parameter space with the reach of the experimental searches, from 1988 to 2021, primarily focusing on collider

experiments. The results here are shown for m_S or m_ϕ below 5 GeV and going down to 0.05 GeV. Across all masses, a mixing angle down to $\sim 10^{-4}$ has been probed. Focusing on the recent results, searches for scalar bosons in final-states of two muons [55, 56] at the LHCb experiment already provide a large coverage in the mass range of $0.3 \leq m_\phi \leq 2$ where the hadronic channel dominates. This is complemented by searches at Belle II [57] where the $ee, \mu\mu, \pi\pi, KK$ final-states have also been considered.

Light BSM particles in exotic Higgs decays

A few searches for light BSM particles with masses less than 5 GeV are presented in this section, targeting their production through Higgs boson decays (i.e. $H \rightarrow SS$). The possibility of such decays is well-motivated since the current measurements of the Higgs boson properties still allow for $\mathcal{O}(10\%)$ deviations from the SM expectations. The recent review and combinations of the Higgs bosons measurements by the CMS experiment [58] further support the possibility of scenarios where the Higgs boson decays into new BSM particles. Furthermore, direct searches for the Higgs boson branching fraction to invisible particles [59] are constrained to below 16%, which is larger than the SM expectation of $\sim 0.1\%$ coming from $H \rightarrow ZZ^* \rightarrow 4\nu$ processes [60]. A global fit of the SM Higgs boson couplings, with the possibility of the Higgs boson to decay to a potentially visible but undetected final state, allows at 95% confidence level (CL) for up to 16% of the Higgs bosons to decay in such a yet undetected way [58].

Search for long-lived particles decaying in CMS muon detectors [61]:

This search performed at the CMS experiment targets BSM particles through their decays to a wide range of final-states, including $ee, \gamma\gamma, KK, \pi\pi$, by identifying their signatures in the CMS muon systems. The analysis is sensitive to particles with large lifetimes. Figure 1.6 (left) shows the results for the $\pi\pi$ channel in the two-dimensional parameter space of the particle mass and the proper lifetime. The analysis is most sensitive to particles with $1 \lesssim c\tau \lesssim 1000$ mm. Note that the branching fraction of $S \rightarrow \pi^+\pi^-$ is assumed to be 100% to reduce the dependency on any underlying model.

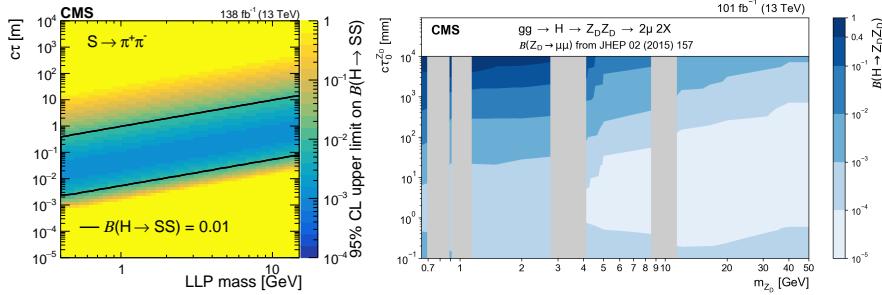


Figure 1.6: Upper limits on the branching fraction of exotic Higgs decays to LLPs by some recent searches by the CMS experiment in final-states of charged pions (left) and muons (right). Figures taken from Ref. [61] (left) and Ref. [62] (right).

Search for long-lived particles decaying into muon pairs with CMS [62]:

This search uses a special dataset collected with the CMS experiment that contains low-momentum muons that are used to identify the long-lived particles decaying through the $\mu\mu$ channel. Figure 1.6 (right) shows the parameter space covered by the search. The analysis is sensitive to lower lifetimes ($0.1 \lesssim c\tau \lesssim 100$ mm) compared to the analysis presented before. Note that several mass regions have been excluded due to the presence of large SM backgrounds. The branching fraction of the particle to muons is taken according to the dark photon model which, at mass=2 GeV, has the same branching fraction as the BSM theory discussed before.

Fig. 1.7 shows the combination of these two CMS results for a few mass points and the corresponding combined coverage in lifetime. A complementary search for a pair of new bosons with masses down to 0.21 GeV in events with four muons in the final-state was also performed in Ref. [64]. However, it is harder to compare the results with Fig. 1.7 due to differences in the theoretical models and the final interpretation. The analysis presented in Chapter 5 aims at providing complementarity in this parameter space by increasing sensitivity to small lifetimes ($c\tau \lesssim 100$ mm) at masses below 2 GeV.

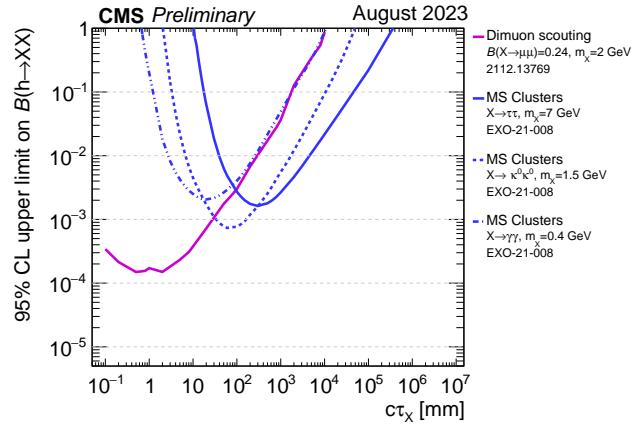


Figure 1.7: Upper limits on the branching fraction of exotic Higgs decays to LLPs as a function of the lifetime for a few selected mass points from various CMS analyses [61, 62]. Figure taken from Ref. [63].

Chapter 2

The CMS experiment at the Large Hadron Collider

This chapter provides an overview of the experimental setup relevant to the data collected and analyzed in this thesis. The material presented here is based on information compiled from various sources.

The history of particle physics can perhaps be traced back to the 19th century, with the discovery of the electron in the cathode-ray experiments performed by J.J.Thomson. Numerous experiments since then have contributed in building the picture of elementary particle physics as we know it today. Over the years, collider experiments have progressively pushed the energy frontier, providing controlled environments to study the Standard Model and to search for new physics.

The Large Hadron Collider (LHC) [65] is the largest and most powerful particle accelerator that produces collisions at the TeV energy scale. In its first 3 years of operation, the LHC was successful in providing the collisions that led to the discovery of the Higgs boson. The LHC has since continued to deliver collisions at center-of-mass energies up to 13.6 TeV.

To record and study these collisions, several large detectors are installed at the LHC, each targeting a rich and broad physics program. The work performed in Chapter 4 and Chapter 5 uses the data collected in one of these detectors, the

Compact Muon Solenoid (CMS) experiment. The following sections provide an overview of the LHC and the CMS detector.

2.1 The Large Hadron Collider

The LHC is located at the European Organization for Nuclear Research (CERN) in Geneva, on the France-Switzerland border. It is the successor to the Large Electron Positron (LEP) collider, using its 27 km circumference tunnel to provide particle collisions at the TeV energy scale. Built between 1998 and 2008, the LHC is a synchotron-type accelerator, performing proton-proton (pp) and heavy ion collisions such as lead-lead (Pb-Pb) collisions [66]. The resulting data is recorded and studied at the ALICE [67], ATLAS [68], CMS [69] and LHCb [70] experiments, as well as several smaller experiments, pursuing distinct physics goals. Since the primary focus of the CMS experiment has been pp collisions, the description of the LHC presented here is centered on that. The information summarized in this section is taken from Refs. [71–73].

The proton-proton collisions occur in the LHC ring after passing through several accelerators at CERN, each successively boosting the center-of-mass energy of the system. A schematic of the various accelerator facilities at CERN is shown in Fig. 2.1. It starts from hydrogen gas, which is first turned into H^- ions. The ions are then guided towards the linear accelerator, LINAC4, which accelerates the beams to an energy of about 160 MeV. The LINAC4 has several components based on radiofrequency (RF) cavities to achieve the gain in energy through oscillating magnetic fields at a frequency of 352.2 MHz. The injection of the ions into these cavities are synchronized with the oscillation, ensuring that all ions are accelerated or decelerated to the desired energy. The synchronized injection naturally leads to successive packs of ions being produced, which are termed ‘bunches’. The ions are then stripped of their two electrons to form proton bunches during the injection into the circular accelerator, the Proton Synchotron Booster.

Within a circular accelerator, two processes occur simultaneously. The first is the increase in energy with every revolution as the beams pass through RF cavities. However, as the energy increases, a stronger magnetic field is required

to maintain the same circular trajectory. This is performed by synchronizing a gradually increasing magnetic field. Based on these principles, the Booster accelerates the beams to 2 GeV and injects it to the Proton Synchotron which then increases the energy to 26 GeV. This is followed by the Super Proton Synchotron (SPS) which accelerates the beams to 450 GeV before injection into the LHC rings.

The LHC ring consists of two counter-rotating beams of protons, circulating in separate beam pipes and merging into a single beam pipe at the interaction points. Operating on the same principles as before, the RF cavities of the LHC ramp up the energies from 450 GeV to half of the center-of-mass energy (\sqrt{s}) for each beam, which is currently $\sqrt{s} = 13.6$ TeV. To bend the beams along the ring, Nb-Ti superconducting magnets maintained at 2 K are operated, reaching a magnetic field above 8 T. Once the beams reach the target energy, the beams are focused with additional magnets at the four interaction points where the four main LHC experiments are situated. At the interaction points, the proton bunches collide at intervals of 25 ns, leading to a rate of about 40 million collisions per second, called the bunch-crossing rate.

Besides the center-of-mass energies, the physics reach of the LHC is also defined by the instantaneous luminosity, \mathcal{L} . For a physics process with cross-section σ , the rate N is calculated as $N = \mathcal{L}\sigma$. Hence, the larger the luminosity, the larger is the rate at which rare processes can be produced at the LHC. Assuming that the beam follows a Gaussian profile, the luminosity is calculated according to Eqn. 2.1.

$$\mathcal{L} = \frac{N_b^2 n_b f_{rev} \gamma}{4\pi \varepsilon_n \beta^*} F \quad (2.1)$$

Here, N_b is the number of protons per bunch, n_b is the number of bunches per beam, f_{rev} is the revolution frequency, γ is the Lorentz factor, ε_n describes the beam spread in position-momentum phase space, β^* relates to the transverse size of the beam at the interaction point and F is luminosity reduction factor depending on the non-zero crossing angle. A detailed description of this is given in Ref. [65].

The amount of data generated at the LHC is measured in terms of the integrated luminosity collected. Since the cross-section is measured in units of

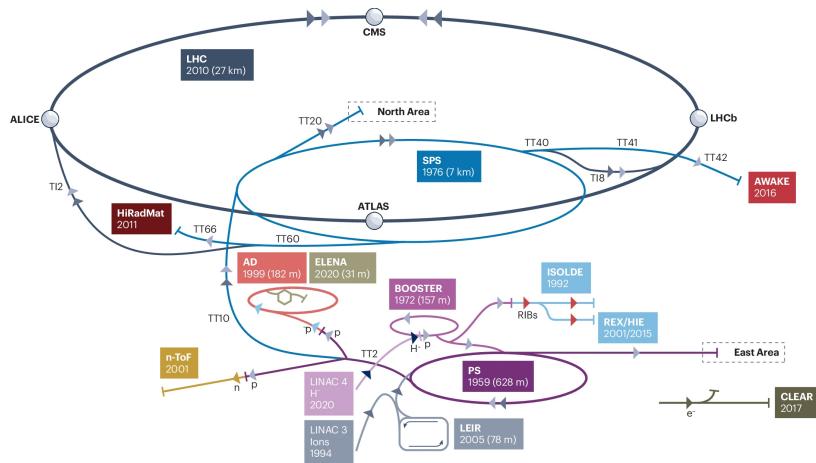


Figure 2.1: A schematic of the CERN accelerator complex. The collision beams pass through the LINAC 4, Booster, PS and SPS before injection into the LHC. Figure taken from Ref. [73].

barn (b), the luminosity is accordingly measured in units of inverse barn or more specifically, femtobarns ($1 \text{ fb}^{-1} = 10^{15} \text{ b}^{-1} = 10^{43} \text{ m}^{-2}$). The LHC has been in operation since 2010, starting with $\sqrt{s} = 7 \text{ TeV}$ and currently, in 2025, running at $\sqrt{s} = 13.6 \text{ TeV}$. So far, the total integrated luminosity recorded by the CMS experiment is 473.84 fb^{-1} at the moment of writing [74].

Based on the operational periods and the long shutdowns (LS) for maintenance and upgrades, there have been three main periods of physics data-taking: Run 1 (2010-2012), Run 2 (2016-2018), Run 3 (2022-2026, planned). This will be followed by LS3, which will be used to prepare for the High-Luminosity LHC (HL-LHC), that has a target of achieving $\sim 3000 \text{ fb}^{-1}$ integrated luminosity.

At each bunch crossing, multiple proton-proton interactions occur simultaneously. Most of these interactions lead to relatively low-energy processes and are termed as pileup. Only a small fraction of collisions result in a hard scattering event with a large momentum transfer since the corresponding cross-section is much smaller than that for soft processes. The resulting particles at the interaction points are detected by the LHC experiments and are used to reconstruct the event.

2.2 Compact Muon Solenoid experiment

The CMS detector [75] is one of the two general-purpose detectors at the LHC, and pursues a broad physics program, ranging from Standard Model physics measurements to new physics searches. Situated at the Point 5 collision point in the LHC ring, it shares similar physics goals as the ATLAS experiment with a distinct detector design.

Since the LHC collisions have a dense hadronic environment at high energies, being produced with each bunch crossing at intervals of 25 ns, there were several considerations that determined the design of the CMS detector. These include radiation-hardness and low occupancy rate (i.e. low number of hits per readout channel to reduce fake rates) achieved through high granularity and good timing resolution. Considering that neutrinos interact only via the weak interaction and are difficult to detect directly, a good momentum resolution and wide fiducial coverage for the reconstruction of the other SM fermions are required to

use the missing transverse energy to account for the neutrinos.

To meet these requirements, the CMS detector, originally proposed in Ref. [76], is composed of several subdetectors, each designed to measure different properties of the particles produced in a collision event. It is 21.6 m long with a diameter of 14.6 m and a total weight of $\sim 14,000$ tonnes. With a cylindrical architecture composed of a barrel part and two endcap parts, the driving feature for the CMS detector is the 3.8 T superconducting solenoid magnet that bends the charged particles and allows for a precise momentum measurement. Within the bore of the solenoid, there is a tracker system and the electromagnetic and hadronic calorimeters, while the muon system is located outside the magnet. A schematic of the CMS detector is shown in Fig. 2.2. The detector components are paired with a trigger system which filters and records the interesting collision events using information from the subdetectors. A description of each of the detector components is provided in the following sections.

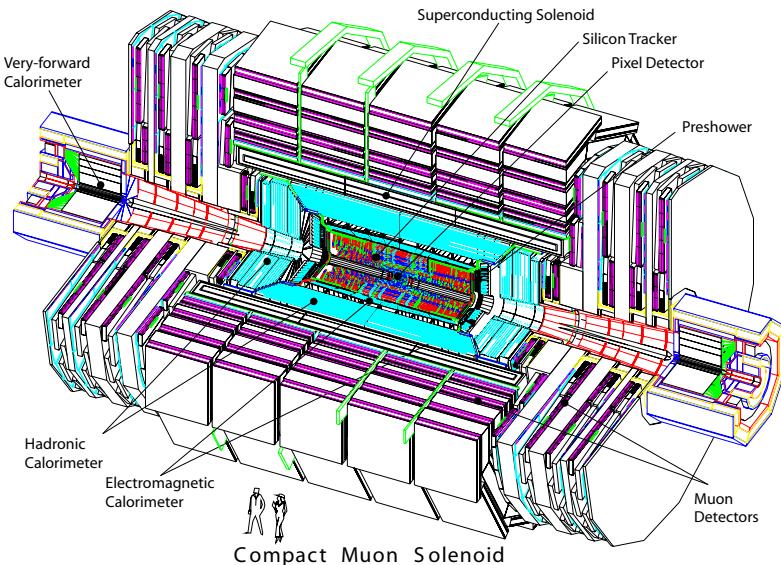


Figure 2.2: A schematic of the CMS detector and its components. Figure taken from Ref. [77].

2.2.1 CMS coordinate system

The CMS experiment uses a standardized reference system, with the collision (or interaction) point at the center of the detector, and is taken as the origin of the coordinate system. The z-axis lies along the beam pipe, while the x-axis points towards the center of the LHC ring and the y-axis is towards the surface. Owing to the cylindrical shape of the detector, points on the transverse or xy-plane are often defined by the radial distance r , and the azimuthal angle ϕ , measured as the inclination from the x-axis. The polar angle θ is measured as the inclination from the z-axis. An illustration of the coordinate system is shown in Fig. 2.3.

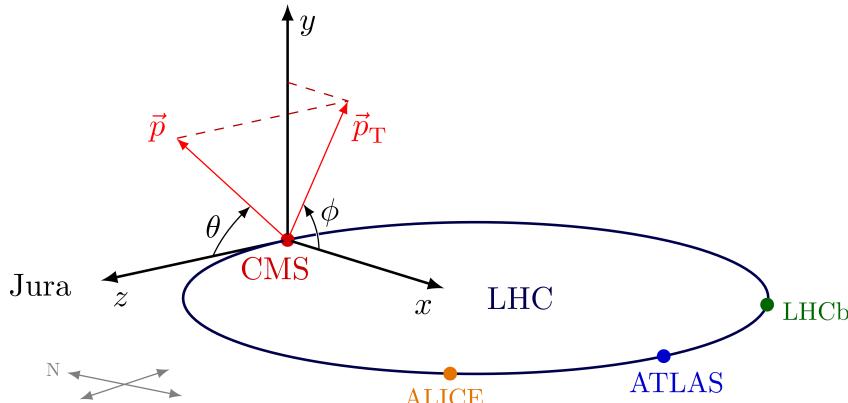


Figure 2.3: A schematic of the coordinate system. Figure taken from Ref. [78].

In the transverse plane, the total momentum of the particles in the event is expected to be zero, since the colliding beams travel along the z-axis with no momentum in the transverse direction. Therefore, quantities measured in the transverse plane are lorentz invariant to boosts along the z-axis. Accordingly, the transverse momentum, p_T , is a commonly used variable at the CMS experiment.

With similar reasoning, a quantity called the pseudorapidity, η (defined in Eqn. 2.2), is often used to describe the inclination from the z-axis.

$$\eta = -\ln \left(\tan \frac{\theta}{2} \right) \quad (2.2)$$

For ultra-relativistic particles, the pseudorapidity is approximately equal to the rapidity, $y = \frac{1}{2} \ln \left(\frac{E+p_z}{E-p_z} \right)$. Since the difference in rapidity is Lorentz invariant

under boosts along the z -axis, the corresponding differences in pseudorapidity, $\Delta\eta$ (under the ultra-relativistic limit), is also lorentz invariant. Therefore, η is preferred over the polar angle θ to describe the inclination from the z -axis.

2.2.2 Silicon tracker

The CMS silicon tracker [79, 80] is the innermost detector surrounding the beam pipe, which has an outer radius of about 22 mm. The first tracker layer is positioned at a radial distance of 29 mm from the z -axis. The tracker measures the three-dimensional position of charged particles at multiple points along their path as they traverse through it. This information is used to reconstruct the helical path of the charged particles due to the magnetic field generated by the CMS solenoid magnet. The reconstructed trajectories in turn allow the determination of the charged particle's momentum. Furthermore, intersections of the reconstructed trajectories enable vertex reconstruction, which in turn allows identification of the primary hard-scattering vertex and pileup vertices.

The CMS tracker system is entirely composed of silicon-based technology. This choice is driven by the following requirements for the tracker – (1) radiation hardness as the proximity to the interaction point leads to intense particle flux, (2) high granularity since every collision can result in production of $O(10^3)$ particles, (3) fast response to deal with the frequency of the particle collisions (~ 25 ns), (4) minimal material to limit multiple scattering and energy loss.

The CMS tracker is divided into two main subsystems: the pixel detector, which is closest to the interaction point and the strip detector, which surrounds it. A schematic of the tracker is shown in Fig. 2.4 where the pixel detector layers are shown in green and the remaining outer layers belong to the strip detector. Together, the tracking system provides coverage up to $|\eta| < 2.5$. While the geometry for the two subsystems varies, they both use the same underlying detection principle. The tracker sensors consist of a p-n type sensor operated in reverse-bias mode, producing a large depletion region. When a charged particle passes through this region, the ionization produced leads to a pulse that is processed by a readout chip to amplify and generate an electronic signal. Multiple pixels or strips are clustered together to form a ‘hit’ which is later used for track reconstruction. As the tracker faces the most radiation, the tracker is operated at

-20° C degrees to reduce the noise in the electronics due to thermal fluctuations.

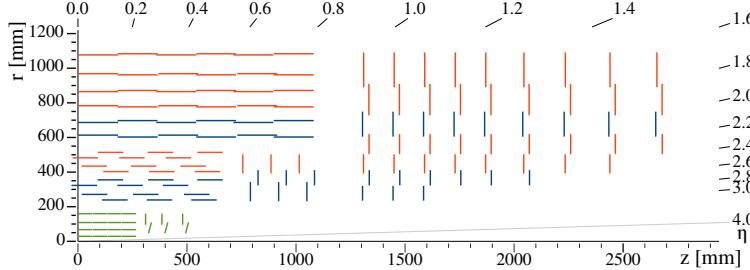


Figure 2.4: Schematic of one quarter of the Phase-1 CMS tracking system. The pixel detector is shown in green, while the strip detector is shown in red and blue. Figure taken from Ref. [81].

The pixel detector consists of four layers in the barrel region and three layers in each of the endcap disks. These pixel layers consist of nearly 124 million pixels, each with a size of $150 \times 100 \mu\text{m}$. The original design consisted of three barrel layers and two layers in each of the endcap disks. At the start of 2017 data-taking, the extra layers were added and the configuration was changed to its current design to deal with the increased luminosity in 2017. This allows for a four-hit reconstruction and provides increased accuracy of the tracking algorithm since the first layer is closer to the beam pipe in the current configuration.

The strip detectors face a relatively reduced particle flux and accordingly, the need for high granularity is lower. The design of the detector consists of several layers of silicon micro-strip detectors, which have a typical cell size of $10 \text{ cm} \times 80 \mu\text{m}$ in the first few layers, increasing in size up to $25 \text{ cm} \times 180 \mu\text{m}$ in the outer region of the tracker. The readout and amplification is performed by the APV25 chip. In early 2016 data, a saturation effect in the APV25 chip was observed under the new data-taking conditions of Run 2, leading to a loss of track reconstruction efficiency due to missing hits [82]. This effect was mitigated in the late 2016 to 2018 data-taking conditions. It will be taken into account when defining the data samples for the analysis in Chapter 5.

The tracker has excellent hit position resolution. For the pixel detector, the transverse coordinate hit resolution is $\sim 10 \mu\text{m}$ and the longitudinal coordinate has a resolution of $20 - 40 \mu\text{m}$ [80]. For the strip detector, the spatial resolution ranges from 20 to $40 \mu\text{m}$ [83].

2.2.3 Electromagnetic calorimeter

The electromagnetic calorimeter (ECAL) [84] is the next radially outward sub-detector which is responsible for measuring the energy of electrons and photons. Upon incidence of an electron or photon on the ECAL, Bremsstrahlung and pair production processes induce electromagnetic showers as they pass through the detector by interacting with the material.

The energy deposited through the showers depends on the radiation length which is inversely proportional to the atomic number of the material. By precisely measuring the total energy of these showers, the energy of the originating electron or photon is determined by the ECAL. This requires the ECAL to be sufficiently long to contain the full shower while also ensuring a good granularity to contain the transverse extent of the shower (defined by the Moliere radius).

Keeping these design requirements in mind, the ECAL is constructed as a homogeneous calorimeter with PbWO_4 scintillating crystal cells. PbWO_4 has a radiation length of ~ 0.89 cm and a Moliere radius of 2.2 cm, allowing for a compact design. The light produced through the scintillation during the electromagnetic shower is read out by photodiodes. The light yield is captured by it and in turn converted to a signal that is proportional to the energy deposit of the incident particle.

The ECAL is divided into two regions, the barrel ECAL (EB) in $|\eta| < 1.479$ and the endcap ECAL (EE) in $1.479 < |\eta| < 3.0$. The EB contains 61200 crystals, each with a front face of $22 \times 22 \text{ mm}^2$ and $26 \times 26 \text{ mm}^2$ at the rear, and provides a granularity of $\Delta\eta \times \Delta\phi = 0.0174 \times 1^\circ$. For the EE, the crystal size is slightly larger comparatively with a front face of $28.6 \times 28.6 \text{ mm}^2$. In addition to these components, a preshower detector (ES) is installed in the region $1.536 < |\eta| < 2.6$ in front of the EE to allow for identification of photon pairs coming from π^0 decays. The energy resolution is studied using all electrons from $Z \rightarrow e^+e^-$ events and is found to be 2–5% across the different years of data-taking in Run 2 [85], as shown in Fig. 2.5.

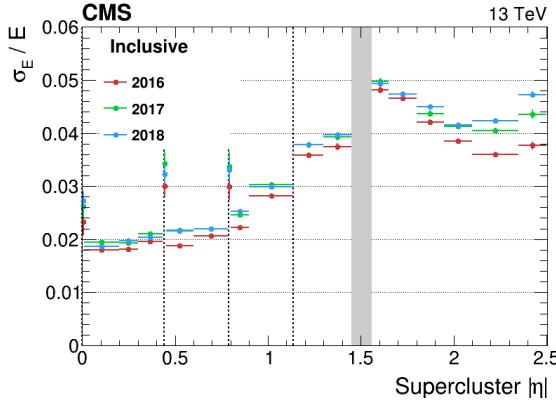


Figure 2.5: Relative energy resolution across different eras of Run 2 in bins of supercluster $|\eta|$. The supercluster corresponds to the clustered group of cells that capture the energy deposition of the electron. Figure taken from Ref. [85].

2.2.4 Hadronic calorimeter

The Hadronic calorimeter (HCAL) [86] lies after the ECAL and captures the energy of the hadrons (charged and neutral). A hadronic shower has a complicated pattern with cascades of multiple types of particles produced as a result of the interaction with the nuclei. The extent of the shower is determined by the interaction length of the absorber material.

The HCAL is designed as a sampling calorimeter, containing alternating layers of absorber and active materials. In this setup, the absorber layer primarily consists of a heavy material such as brass, to trigger the hadronic showers, while the active layer, consisting of plastic scintillators, measures the energy of the shower. The lost energy in the absorber layers is accounted for through calibration and corrections for this effect.

A schematic of the HCAL is shown in Fig. 2.6. The HCAL consists of 4 parts: HCAL barrel (HB) in $|\eta| < 1.3$, HCAL endcap (HE) in $1.3 < |\eta| < 3$, HCAL outer (HO) in $|\eta| < 1.3$, and HCAL forward (HF) in $3 < |\eta| < 5.2$. Since the HB is radially restricted by the inner extent of the magnet (discussed later), the thickness is insufficient to absorb the highest-energy hadronic showers. The HO, placed outside the magnet in the barrel region, compensates for the missing sampling depth. The HB, HE and HO have plastic scintillators as the active

material which use similar principles as the ECAL to convert the light yield to a measurement of the deposited energy. The HF uses quartz fibres as active material with the signal generation based on the Cherenkov radiation principle. Brass mostly forms the absorber material for the HB and HE while the HO uses the solenoid as an absorber layer in most parts. Besides this, steel is used in some regions of HB and HO and for the HF.

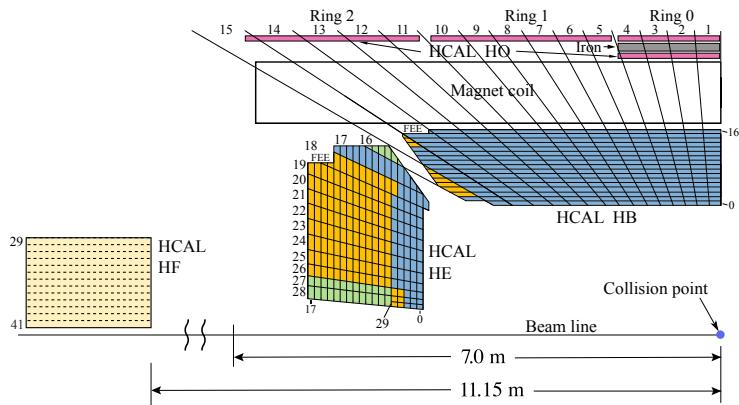


Figure 2.6: Schematic of one quarter of the HCAL system. Figure taken from Ref. [87].

The HCAL is segmented in $\eta - \phi$ units called towers with a granularity of $\Delta\eta \times \Delta\phi = 0.087 \times 5^\circ$ for $|\eta| < 1.6$ and $\Delta\eta \times \Delta\phi = 0.175 \times 10^\circ$ for $|\eta| > 1.6$, with further segmentation in depth. Since hadrons are typically produced as a collimated spray of particles called jets (discussed in the next chapter), the performance of the HCAL is evaluated through the jet energy resolution, which takes inputs from the entire CMS detector.

2.2.5 Solenoid Magnet

The presence of a magnetic field causes charged particles to follow a helical path which enables the determination of two important quantities. One is the sign of the charge of the particle while the second is the measurement of the particle momentum. For a precise measurement of the momentum to occur within the volume of the CMS detector, a strong and uniform magnetic field is required

which ensures a sufficient curvature in the path of a particle with energies up to $\mathcal{O}(1000\text{GeV})$.

With this design requirement, the magnet at the CMS detector [75] is a superconducting solenoid magnet which is 12.9 m long in the z-axis with an inner diameter of 6.3 m. It is placed between the hadronic calorimeter and the muon systems, covering the barrel region of the detector. The solenoid coils are made from niobium-titanium conductors and are cooled to 4 K to generate a uniform field of 3.8 T. The magnet is surrounded by iron return yoke on the outside. The return yoke, weighing 10,000 ton, serves a dual purpose of confining the magnetic flux and stopping all particles besides muons and neutrinos.

2.2.6 Muon system

The muon system lies beyond the solenoid magnet and is the outermost layer of the CMS detector. It provides a precise reconstruction and identification of muons, reaching up to 1–3% in p_T resolution for muons with $p_T \leq 100$ GeV. The system consists of gas ionization chambers placed among the return yoke layers, as shown in Fig. 2.7. The ring of chambers assembled there are called muon stations.

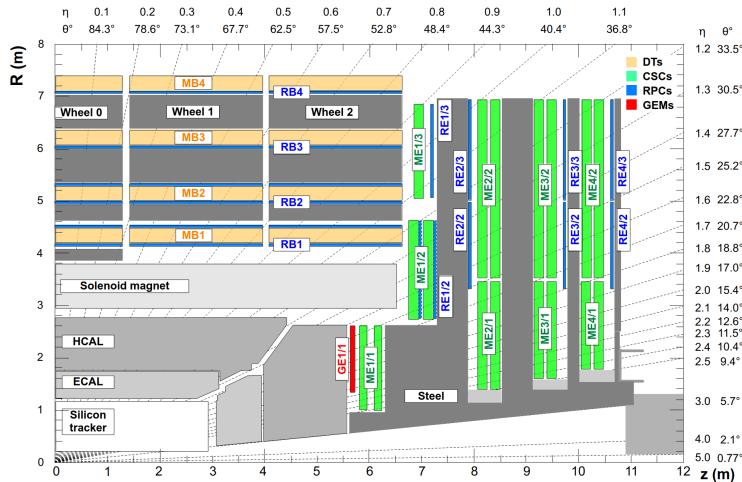


Figure 2.7: A schematic of the muon system shown in one quarter of the CMS detector. Figure taken from Ref. [88].

Three types of chambers have been used for the muon system: drift tube chambers (DTs) in $|\eta| < 1.2$, cathode strip chambers (CSCs) in $0.9 < |\eta| < 2.4$ and resistive plate chambers (RPCs) in $|\eta| < 1.9$.

The DTs consist of multiple layers of staggered drift cells which evaluate the position of the muon by measuring the drift time of the electrons to the central anode wire. The layers are grouped into superlayers with alternating orientations which respectively measure the z coordinate and the position in the $r - \phi$ plane. This provides an excellent spatial resolution.

The CSCs, in the endcap region, consist of anode wires crossed with cathode strips where the wires measure the r coordinate while the strips provide the ϕ coordinate. The depth indicates the z coordinate.

In $|\eta| < 1.9$, the DT and CSC modules are supplemented by the RPCs. They are double-gap chambers which primarily provide a precise timing information. This is useful in assigning the muons to their respective bunch-crossings. Since Run 2, gaseous electron multiplier (GEM) detectors have been added in $1.6 < |\eta| < 2.2$, in front of the ME 1/1 muon station, to improve the measurements of the muon trajectory.

2.2.7 Trigger system

The pp collisions at the LHC occur at a rate of 40 MHz, as mentioned in the previous section. Storing all the information of an event amounts to an average size of 1 MB. At the rate of LHC collisions, this would translate to 4 TB/s of data being generated at the CMS experiment and is not practically feasible to store. When comparing the cross-sections of various processes, the cross-section of interesting processes such as the production of Z bosons, is many magnitudes smaller than that of the dominant low energy multi-jet processes. This further provides the motivation for a dedicated selection which stores only the events with potentially interesting physics in them. This filter system is called the trigger system. The CMS experiment employs a two-tier trigger system [89], the first one being the Level-1 hardware trigger (L1 trigger) followed by the software-based High-Level trigger (HLT). The L1 trigger reduces the rate from 40 MHz to 100 kHz followed by the HLT which further reduces the rate to ~ 1 kHz.

The L1 trigger utilizes custom-designed electronics to enable fast trigger de-

cisions within the maximum allowed latency of $4\mu\text{s}$. In the time period between two bunch crossings, the L1 trigger system executes algorithms to arrive at a decision based on inputs from various subdetector systems. As a result of the time limitation, the L1 trigger system does not use information from the tracker. The L1 trigger system uses low-level information such as hits from the muon systems and energy deposits from the calorimeters, to construct regional objects called trigger primitives. They are then subsequently passed to other algorithms which construct a crude image of the event in terms of the number of physics objects such as muons or electrons. Eventually, the information is globally considered while taking the final L1 trigger decision. During the time needed to reach a L1 decision, the data is stored in memory buffers. An overview schematic of the L1 trigger architecture is shown in Fig. 2.8.

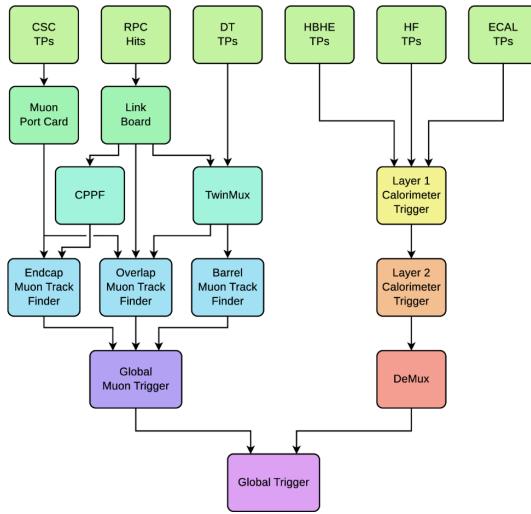


Figure 2.8: Schematic of L1 trigger workflow. Figure taken from Ref. [90].

The next stage in the triggering process is the HLT which uses the full event information (including tracker) to select or reject events. The HLT is a large computing farm consisting of CPUs and GPUs that performs a lighter version of the final (offline) reconstruction algorithms. As a result, a refined picture of the event with well-reconstructed physics objects is already available at the HLT where complex selections targeting specific physics analyses can be imple-

mented. The workflow described here is for the standard data stream used at the CMS experiment. There are alternate data-taking techniques that were also used during Run 2, such as data scouting and data parking [91].

Chapter 3

Simulation and reconstruction of physics objects

This chapter presents an overview of event simulation tools and physics object reconstruction techniques used at the CMS experiment. This chapter builds upon the previous chapter, providing the link between the detector signals and the final reconstructed event. The discussion here is based on a combination of sources.

For any experiment, an important aspect is to have a good understanding of the underlying theory and how it reflects the observations of the experiment. For the proton-proton collisions at the CMS experiment, the Standard Model provides the theoretical framework within which the particle interactions take place. However, there is still a need for the link between the underlying theory and what we see in the experiment. Due to its complex nature, a detailed modelling of the collision events is required to not only confirm the expectations of the Standard Model but also to guide the searches for BSM physics. This modelling, called the event simulation, is followed by detector simulation, to describe how the resulting particles interact with the CMS detector. Together, they form the first step of any physics analysis and provide the basis for identifying the key variables relevant to it. The details of event simulation and detector simulation are presented in Section 3.1.

The previous chapter described how the particles interact with the detec-

tor components to produce digital signals. However, the raw detector response from these interactions is not sufficient to detect the particle and reconstruct its properties. To achieve this, the information from various parts of the detector is combined with dedicated reconstruction algorithms. They transform the low-level detector signals into high-level physics objects, such as muons, electrons, photons, jets and the missing energy, that provide a complete description of the event. The details of the physics object reconstruction in CMS are discussed in Section 5.4.1.

3.1 Event simulation

Simulating pp collisions is a highly challenging and complex task. Starting from the matrix element calculation for a pp collision, the relevant computation of the process can become difficult to compute beyond a few orders. Depending on the physics considered, a converging perturbative expansion may not even be possible, such as in the case of non-perturbative QCD. This is followed by integrating over the multi-particle phase space, taking into account that intermediate states can radiate or decay into additional particles. While these tasks are performed by simulations, performing them analytically still remains extremely difficult. As an alternative, simulation software packages, known as event generators, use Monte Carlo techniques and approximate models to simulate the collision events. A comprehensive review of event generators can be found in Ref. [92].

The first part of the simulation is the generation of the hard-scattering process which can be defined as the process which has a large momentum transfer. In our events of interest, since by far most pp collisions happen at low momentum transfer, typically only one occurs with a large energy transfer. Since protons are composite particles containing quarks and gluons (partons), the hard-scattering process at the energy scale we consider takes place by the interaction of two partons.

The cross-section for a hard-scattering process can be computed with collinear factorization [93]. In this formulation, the cross-section is factorized into two components, the perturbative part and the non-perturbative part. The perturbative part is evaluated to a particular fixed order by calculating the parton-level

cross-section of the hard scattering process. The non-perturbative part is determined with the Parton Distribution Functions (PDFs) that are the probability densities of a parton carrying a certain fraction of the proton's momentum.

The parton-level cross-section is calculated from the matrix element squared and its integral over the kinematic phase space, while also taking into account the possibility of different spin-states of the initial states and colour degrees of freedom. This is typically performed by evaluating Feynman diagrams for the process at a fixed order through matrix element generators such as POWHEG [94–96] (next-to-leading order accuracy), Madgraph5 [97, 98] (leading order accuracy) and Madgraph5_aMC@NLO [99, 100] (next-to-leading order accuracy) which provide an automated way to compute these. At a particular factorization scale, the PDFs are evaluated according to various schemes, such as those provided by NNPDF [101]. This completes the calculations needed for producing the hard-scattering process.

There is still a possibility for the incoming or outgoing particles to radiate more particles. For instance, before the hard-scattering, the quarks can radiate gluons which in turn can split into more partons. This is called initial-state radiation (ISR). When this radiation occurs after the hard-scattering process from the outgoing partons, it is termed as final-state radiation (FSR). These radiations or emissions are collectively described by parton showers. Exact calculations for the evolution of the partons and all the successive emissions are computationally expensive. The parton shower framework combines the fixed order description based on exact calculations that are applicable where the perturbation theory is valid (i.e. above $\Lambda_{\text{QCD}} \sim 300$ MeV), with approximate modelling of the successive emissions. In the non-perturbative regime, the evolution of the particles is performed by phenomenological models such as the Lund string model [102], until colour-neutral final-states are reached. This process is called hadronization.

The above parton shower and hadronization processes are typically carried out together by event generators such as PYTHIA8 [103]. Furthermore, there is also a possibility for other partons in the proton to interact. This typically results in low-energy activity in the event and is termed as underlying event. PYTHIA8 has several models to simulate the underlying event such as Multi-parton interactions (MPI) [104]. The decays of short-lived particles is also simulated within

the event generators. A sketch of an event as viewed by an event generator is shown in Fig. 3.1.

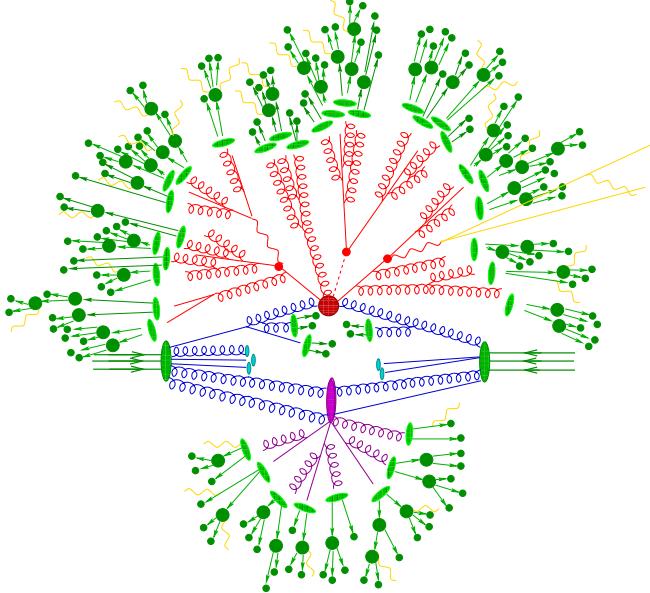


Figure 3.1: An illustration of a proton-proton collision as simulated by a event generator. The red blob in the center represents the hard scattering. The red lines represent the perturbative parton shower. The two big light green blobs are the incoming protons. The purple blob shows a secondary parton interaction. The remaining green lines and blobs represent the hadronization processes and resulting hadrons. The yellow indicates photon radiations. Figure taken from Ref. [105].

Finally, the event details are passed over to the detector simulation. The CMS detector is modelled in Geant4 [106], which is a software that simulates the interaction of particles as they pass through materials. During a single bunch crossing, multiple pp interactions occur at low-energies, resulting in pileup. PYTHIA8 is also often used to perform the simulation of these events (termed as minimum bias events) which are then overlaid on the hard-scattering process.

3.2 Physics object reconstruction

As the particles travel through the detector, their signatures are captured as signals in the various components. This low-level information is progressively combined into higher-level objects. In the tracker, individual hits are clustered to form track segments which eventually form track candidates for charged particles. Similarly, track segments for muon candidates are reconstructed from hits in the muon chambers. In the calorimeters, energy deposits in the cells are grouped into clusters which subsequently provide the information to reconstruct jets, electrons, photons and hadrons. The information from all subdetectors is supplied to a central algorithm (the particle flow algorithm [107]) which eventually reconstructs the physics objects that are used in analyses. The main set of physics objects includes muons, electrons, photons, jets, missing energy, among others. In the following sections, some commonly used physics objects are described.

3.2.1 Tracking

The CMS tracking system is responsible for reconstructing the tracks of charged particles produced in pp collisions. As a charged particle passes through the tracker, it interacts with the modules to produce hits. The tracking algorithm is then tasked with the duty of reconstructing the particle trajectory which follows a helical path as result of the magnetic field. With the design projection of roughly 1000 charged particles per bunch crossing at $\mathcal{L} \sim 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ [108], the busy environment of a pp collision puts demanding requirements on the tracking algorithm to attain a high track reconstruction efficiency while maintaining a low fake rate. Some additional challenges include reconstruction of tracks over a broad range of p_T , separation of nearby tracks and maintaining excellent impact parameter resolution.

With these targets in mind, the CMS experiment employs an iterative tracking algorithm based on the Kalman filter method [109–111]. The main philosophy behind an iterative approach is to reduce the combinatorial complexity of the large number of possible hit combinations. This is achieved by reconstructing the simpler tracks (such as high momentum tracks) in the initial iterations and

masking the associated hits. The algorithm then progressively moves towards more complex tracks but with a reduced number of possible combination of hits due to the masking from previous iterations. As a result, each iteration allows for reconstruction of more challenging tracks.

Each iteration consists of four main steps. The first step is the track seeding which involves the creation of a track segment using 2–4 hits. Additional requirements on the p_T and on geometrical parameters, like the transverse distance relative to the beam spot, are placed on the track seed. Each iteration uses a different type of track seed. The initial iteration uses seeds with four pixel hits with relatively high p_T before gradually relaxing the requirements for the track seed parameters and allowing a smaller number of hits in later iterations. In the later iterations, seeds include some or all of the hits from the strip detector.

The second step is the track finding/building step based on the Kalman filter technique. In this step, the compatible hits in detector layers beyond the last hit of the track seed are successively added to extend the seed trajectory. At each layer, the track parameters are updated, eventually building track candidates as the outermost layers of the tracker are reached. These parameters define the equation of motion for the trajectory which, in turn, determines the hit compatibility based on the location of the hit. The average energy loss of the particle is also accounted for during this track extrapolation. If no compatible hits are found in the layer, a ghost hit is included according to the trajectory. This accounts for any lost hits that may occur due to reconstruction inefficiency or non-functioning detector components. In case of more than one compatible hits, new trajectories are created and are then propagated independently. Once the tracks are completed, a backward track fitting is performed, starting with the outermost hit. This leads to the third step, called the track fitting step, which refits the trajectory using the inside-out and the backward track fitting results. This improves the precision of the track parameters and deals with any bad hits in the trajectory.

The final step consists of the track selection. This is an essential step for suppressing the number of fake tracks in the event. In this step, the tracks are classified into different categories based on their quality. A machine learning based classifier is employed at this stage to select the high-purity tracks.

The details of the track finding and track selection algorithms for Run 2

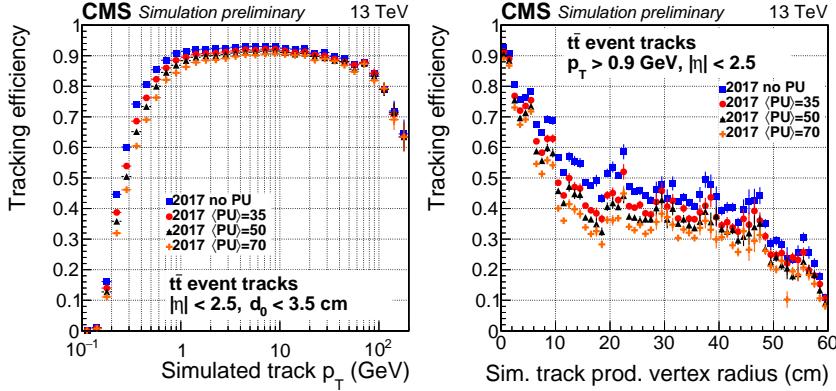


Figure 3.2: The track reconstruction efficiency as a function of simulation track p_T (left) and simulation track production vertex radius (right) at different pileup conditions. Figures taken from Ref. [115].

and Run 3 are mentioned in Refs. [112] and Refs. [113, 114] respectively. Focusing on the performance in Run 2, the tracking efficiency in simulation has been explored in Ref. [115] using $t\bar{t}$. Fig. 3.2 shows the tracking efficiency as a function of the track p_T (left) and the track production vertex (right). For the nominal pileup of $\langle \text{PU} \rangle = 35$ in 2017, a track reconstruction efficiency of $\sim 90\%$ is present for $p_T > 1 \text{ GeV}$. The efficiency drops as the tracks are produced further away from the beamspot. Studies in Ref. [116] compare the tracking efficiency using $Z \rightarrow \mu\mu$ events in simulation and data, and found them to be in agreement within 1–2%. A complementary measurement using charged pions produced from D-meson decays [117] shows that the tracking efficiency in simulation and data is consistent within 2–3% uncertainty.

In the context of the search for light scalars decaying to charged hadrons or muons, the tracking performance has a direct impact on the reconstruction of the final-state particle momenta and position. The high tracking efficiency for tracks with p_T as low as 1 GeV allows for the inclusion of low- p_T tracks in the analysis. Conversely, the drop in tracking efficiency at large displacements highlights the challenging aspect of reconstructing the decay products of long-lived scalars.

3.2.2 Primary vertex and vertex reconstruction

When colliding proton bunches overlap, several pp interactions can take place, producing several interaction points, called vertices. Each interaction point produces particles that are then reconstructed in the CMS detector. In Run 2, the average number of interactions per bunch crossing ranged from $\langle \text{PU} \rangle = 27 - 38$ [74]. However, most of these interactions are low-energy collisions that typically do not involve interesting physics processes. Accordingly, these vertices are termed as pileup vertices. In contrast, the hard-scattering vertex, termed as the primary vertex, is usually the interaction of interest and leads to the final-state particles that are studied in physics analyses. It is vital to reconstruct each of these vertices to identify the particles of interest and also to reject particles coming from pileup which can contaminate the event.

The vertex reconstruction is performed with the adaptive vertex fitter [118] which is based on the Kalman filter method [109–111]. The algorithm starts with an initial vertex candidate using it as a seed position based on the intersection of tracks and constrained by the beam line. The associated tracks are then assigned weights based on their compatibility with the candidate vertex. In particular, the impact parameter of the tracks is calculated with respect to the candidate vertex, down-weighting the tracks which lie further away. The algorithm iteratively minimizes the weighted squared distances of the tracks from the vertex and updates the track weights until the fit converges. This procedure is performed for all vertices obtained from the set of initial vertex candidates. The vertex with the largest sum of track p_T^2 is considered the primary vertex. The pileup vertices are expected to be very close to the primary vertex in the transverse plane (or xy -plane) but spread up to few cm along the beam line (or z -axis). In Run 2, the vertexing resolution in the transverse and longitudinal directions is $10 - 20 \mu\text{m}$ for vertices with more than 60 tracks [115].

In addition to the primary and the pileup vertices, there is also the possibility of vertices that are displaced in the transverse plane. These vertices are referred to as the secondary vertices and arise from the decays of hadrons which travel a macroscopic distance before decaying. Reconstruction of secondary vertices is an essential component in identifying b -quark jets in the events. It also plays

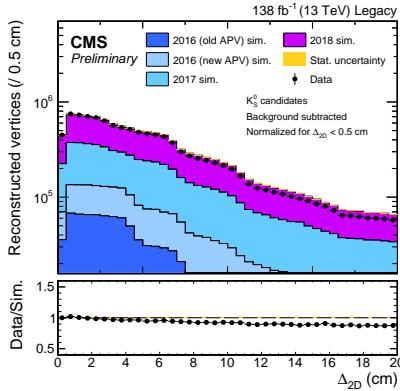


Figure 3.3: The number of reconstructed K-meson vertices as a function of the vertex displacement in simulation and data for Run 2. Figure taken from Ref. [119].

an important role in several long-lived searches for new physics. Specifically, for the search for light scalars decaying to charged hadrons or muons, the reconstructed secondary vertex provides the necessary information to discriminate signal from background processes. Additionally, the track parameters can be improved in precision using the fitted vertex position as a constraint. A relevant study on the performance of displaced vertexing in Run 2 is presented in Ref. [119]. The study reconstructs the secondary vertices of the neutral K meson (K_S^0) which have a mean decay length of ~ 27 mm, before decaying into two charged pions. Fig. 3.3 shows the performance in data and simulation for Run 2 and indicates that simulation matches to data to within 10-15% at large displacements. This method is revisited in the context of the search for light scalars presented in Chapter 5.

3.2.3 Particle flow algorithm

Information from the various subdetectors of the CMS experiment is processed by the Particle Flow (PF) algorithm [107] to identify the particles in the final state of the event. The PF algorithm classifies the particles into 5 physics objects, based on their signature through the detector. Figure 3.4 illustrates the unique signature of these physics objects, namely, muons, electrons, photons,

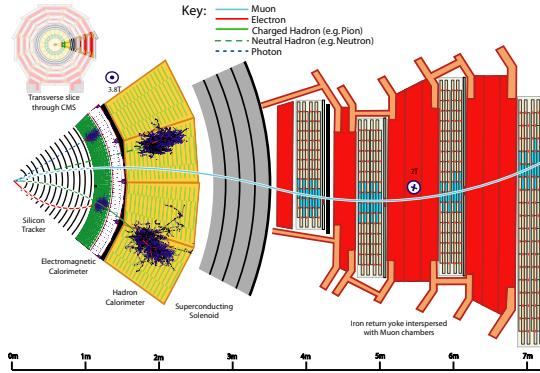


Figure 3.4: Sketch of the physics objects reconstructed by the PF algorithm as they traverse through the CMS detector. Figure taken from Ref. [107].

charged hadrons and neutral hadrons. This list of particles forms the basis for the complete event description by providing information for jet reconstruction, missing transverse energy, particle isolation and more.

The particle flow algorithm follows a linking algorithm where it connects or geometrically matches the tracks in the tracker system to the rest of the detector and successively identifies particles, starting with the muons. The muons provide the cleanest signature in the detector with tracks in the muon chambers and very small energy loss in the calorimeters. The particle flow algorithm reconstructs and classifies them into different muon types based on whether a corresponding track is found in the tracker system. The hits, tracks and energy deposits corresponding to muons are then removed.

With the information from the calorimeters, the remaining tracks are extrapolated outwards, searching for deposits (or energy clusters) in the ECAL and HCAL. Since electrons deposit most of their energy in the ECAL, they are identified as tracks with an associated ECAL cluster only. The remaining ECAL and HCAL cluster energies are compared with the sum p_T of the tracks geometrically matched to the clusters to classify the charged hadrons, neutral hadrons and photons. If the cluster energy is compatible with the sum p_T of the associated tracks, each track is identified as a charged hadron. This, in particular, is important in reconstructing collimated charged hadrons which may have overlapping en-

ergy deposits in the calorimeters but are well reconstructed as individual tracks. Calorimeter clusters with energies larger than the corresponding tracks, indicate that there is also a contribution from neutral particles. This is taken care of by first assigning the tracks as charged hadrons and removing the corresponding energy from the cluster. Based on the fraction of energy between ECAL and HCAL, the remaining energy is reconstructed as a photon (if the majority of energy is in ECAL) or as a neutral hadron. Finally, calorimeter deposits with no associated tracks are similarly assigned as photons or neutral hadrons.

The PF objects reconstructed here are typically further refined with additional identification requirements depending on the needs of the physics analysis. This is described in the following subsections.

3.2.4 Electrons and Photons

The electrons and photons deposit nearly all their energy in the ECAL. Additionally, the electron also leaves hits in the tracker which, in principle, is expected to lead to a distinct track associated with an energy deposit in the ECAL. However, the bremsstrahlung process for electrons and e^+e^- pair production process for photons due to the material before ECAL, can cause an electromagnetic cascade which can lead to a significant energy loss and deviation in the particle's trajectory. Therefore, dedicated algorithms are required to accurately reconstruct the energy of the primary electron or photon. In the previous section, the PF reconstruction procedure for the electrons and photons was briefly outlined. The detailed reconstruction procedure is described below.

Starting with the ECAL, the crystals are grouped together to form ECAL clusters based on the energy deposited by the primary particle. Similar to the tracking algorithm, a seed cluster is defined to form a supercluster by combining it with other surrounding clusters falling within a certain window around it. This is required since the energy deposit can potentially come from an electromagnetic shower instead of just the primary particle, leading to a larger spread in the energy distribution.

In the next step, the trajectory of the electrons are reconstructed using two approaches. The first approach is ECAL-driven and searches for a track seed using the inner tracker hits that is compatible with a supercluster. The second

approach is tracker-driven and iterates over all tracks with $p_T > 2$ GeV that have been reconstructed by the tracking algorithm discussed in Section 3.2.1. The seed of any track compatible with an ECAL cluster is used as a track seed for the electron trajectory reconstruction.

A similar iterative procedure as outlined in Section 3.2.1, is performed with the ECAL-driven and the tracker-driven seeds with the Gaussian-Sum-Filter (GSF) [120] instead of the Kalman Filter. The GSF method provides a better description of the energy loss component in the fit, which is crudely assumed to be a gaussian in the Kalman filter method. For electrons where the bremsstrahlung energy loss has potentially a significant contribution in the trajectory, this is an important modification for the electron track reconstruction.

The electron candidate is then constructed by linking the GSF tracks with the superclusters using multivariate techniques. The photon candidates are reconstructed using the superclusters without any associated GSF track. However, since there is a possibility of $\gamma \rightarrow e^+e^-$, an additional algorithm also takes into account any compatible pair of tracks while constructing the photon candidate. A more detailed description of the electron and photon reconstruction and its performance can be found in Ref. [121].

3.2.5 Muons

Muons are reconstructed using signatures in the muon system and/or the tracker system. The muon system provides a relatively clean environment for track reconstruction, since the CMS solenoid magnet and the calorimeters stop most particles from reaching the muon system. Tracks reconstructed solely from hits in the muon system are referred to as standalone muons. As they are not required to have a matching track in the tracker system, this allows for standalone muons to be excellent probes in long-lived particle searches [122, 123]. Muons originating from within the tracker system are efficiently reconstructed as tracks. Their corresponding trajectories in the tracker system are extrapolated outwards to search for matching segments in the muon system. Muons reconstructed in this way are called tracker muons. The final type of muon is the global muon which applies an outside-in approach to reconstruct the full trajectory. It starts from the standalone muon and searches for a matching track in the tracker system. The

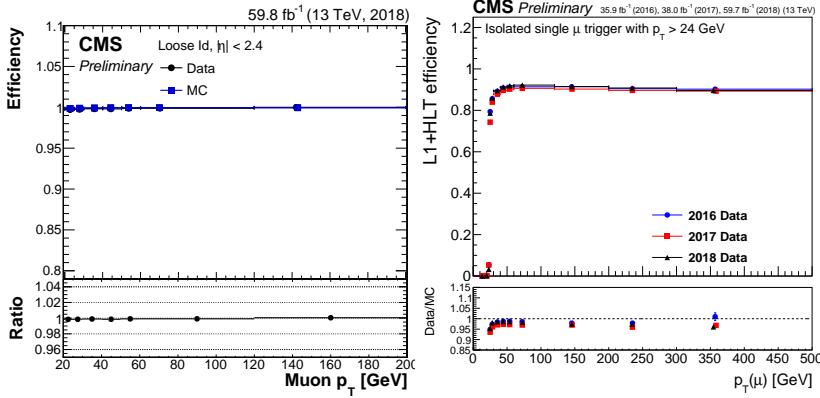


Figure 3.5: The loose muon ID efficiency in 2018 for muons with $p_T > 20$ GeV (left) and the isolated single muon trigger efficiency in full Run 2 (right) as a function of the muon p_T . Figures taken from Refs. [125, 126].

measured properties of the global muon greatly benefit from a combined fit using the tracker muon track and the standalone muon track.

The quality of the reconstructed muon is measured with various identification criteria which are optimized for different physics signatures. This is vital to ensure a high purity of muons and restricting misidentified or poorly reconstructed tracks from being used in physics analyses. Some of the commonly used ones are the loose ID, medium ID and the tight ID [124].

The loose muon ID requires the muon to be either a tracker or a global muon. As the tracker muon and muon segment reconstruction are highly efficient, this ensures that nearly 99% of the muons are reconstructed as global muon or tracker muon or both and pass the loose ID. This requirement is designed to select prompt muons, as well as displaced muons originating from light and heavy flavor decays. The medium and tight identification requirements build on the loose ID requirements by including more constraints on the track quality variables such as the goodness-of-fit. These tighter ID requirements are well-optimized to suppress hadrons misidentified as muons and geared more towards ensuring a high purity of prompt muons.

Muons, like other reconstructed physics objects, often show improved performance in simulation compared to data due to imperfect modeling of detector

inefficiencies present during data-taking. While these effects are typically small, the discrepancies between simulation and data are measured and corrected in simulation. A widely used approach is the Tag-and-Probe method [124], which exploits well-known resonances such as the Z boson and J/ψ , and their decays to muons, to extract efficiencies separately in simulation and data. The resulting data-to-simulation scale factors are then applied as corrections to simulation to provide a more accurate description of object efficiencies. As an example, Fig. 3.5 shows the efficiency of the loose muon ID in 2018 (left) and the isolated single muon trigger efficiency in full Run 2 (right). These measurements are obtained using the $Z \rightarrow \mu\mu$ decays and is generally the preferred choice for studies targeting muons with $p_T > 15$ GeV. For the search for light scalars decaying to hadrons or muons, the isolated single muon trigger is used to select the candidate events. However, the $Z \rightarrow \mu\mu$ topology does not fully represent the kinematics of the targeted signal, and therefore the trigger efficiency measurements shown in Fig. 3.5 do not accurately capture the data-simulation corrections. A dedicated measurement of trigger efficiency measurement for the muon trigger was carried out to obtain a more accurate picture. The details of it are provided in Chapter 5.

3.2.6 PF Isolation

Most of the targeted physics in the CMS experiment aims to capture the particles coming from the hard-scattering process. Muons are often used as probes for such physics processes. While the identification requirements already provide an effective way to increase the purity of signal-like muons, additional selection criteria are needed to provide the signal-to-background discrimination. One such quantity to suppress the background muons is the isolation. This quantifies the amount of hadronic activity around the muon and provides a powerful tool to remove muons originating from jets or unwanted heavy-flavour decays. The isolation can also be calculated and applied to other objects, such as electrons.

Isolation is defined in several ways. One of the popular choices is the PF isolation which uses particle-flow candidates in a cone around the object to measure the relative isolation with respect to its p_T . The cone is constructed using the radius $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$. PF objects falling within $\Delta R < 0.4$ are considered in the calculation of the isolation shown in Eqn. 3.1. The equation is split into two

parts. The first term measures the contribution from the charged hadrons coming from the primary vertex. The second term estimates the neutral contribution using the neutral hadrons and the photons. However, since there is no information on whether the neutral particles originate from the primary vertex, a correction factor $\Delta\beta$ is included to account for contamination from pileup. According to studies in [127], the neutral pileup contribution can be approximated to half of the charged pileup contribution in an event. Therefore, the correction factor is taken to be $\Delta\beta = 0.5\Sigma_{\text{PU}}^{\text{ch. had.}} p_{\text{T}}$.

$$I_{\text{rel}} = \frac{\Sigma_{\text{not PU}}^{\text{ch. had.}} p_{\text{T}} + \max(\Sigma^{\text{neutr. had.}} E_{\text{T}} + \Sigma^{\gamma} E_{\text{T}} - \Delta\beta, 0)}{p_{\text{T}}^{\text{object}}} \quad (3.1)$$

In the context of the light scalars search presented in this thesis, the isolation is an important selection criteria to suppress the background. The PF isolation, as calculated in Eqn. 3.1, is used for the final-state muons. For the hadrons, a slight modification in the PF isolation calculation is introduced to avoid any biases from other targeted final-state objects. This is described in more detail in Section 5.4.7.

3.2.7 Jets

In Section 3.1, the production of partons from proton-proton collisions and their progression towards color-neutral states was discussed. The partons branch into several particles as they progress through the stages of hadronization and fragmentation. The particles then pass through various detector components where they may further radiate or decay into a cascade of particles, leaving their signatures in the detector. Ultimately, a parton is observed as a collimated spray of particles in the detector, referred to as a jet. Jets are primarily composed of hadrons-both charged and neutral-but can also include photons, electrons, and muons resulting from decays of short-lived hadrons. They require efficient clustering algorithms to accurately reconstruct the properties of the originating partons in the event. The clustering is carried out using reconstructed particles in the event to form the detector-level jets (or reconstructed jets). In simulation, the clustering is also performed at the generator-level (i.e. before the final-state

particles pass through detector simulation), to produce particle-level jets.

At the detector-level, the detector response and possible misreconstruction of jet properties require an additional step to translate the measured jet properties to the originating parton properties. The reconstructed jets undergo a set of corrections in this step, termed as the jet energy corrections. This section briefly describes the jet clustering procedure and algorithm(s) adopted in the CMS experiment. The jet energy corrections consist of parts which calibrate the scale (or mean) of the jet energy distribution to what is expected at the particle-level for jets. There is another component which is important to achieve increased compatibility between jets in data and simulation. It refers to the spread in the jet energy distribution compared to the particle-level jet energy and is termed as the jet energy resolution. A detailed study on the jet energy resolution for the Run 2 data-taking period was performed and is presented in Chapter 4.

Jet reconstruction

Experimentally, a detector-level jet can be defined as the collection of final-state particles that most comprehensively represents the properties of the originating parton. It is in particular challenging to establish a direct correlation between the reconstructed jet and the parton since the evolution of a parton into a jet consists of non-perturbative effects like hadronization and underlying event contamination. As a result, a jet reconstruction algorithm (i.e. jet clustering algorithm) needs to be robust against such effects. There are several clustering algorithms (see examples in [128, 129]) available which broadly fall under two categories - sequential recombination and cone algorithms - with the aim to balance computational speed and resilience to QCD effects. In the CMS experiment, the anti- k_t algorithm [130] is a widely used clustering algorithm which is based on sequential recombination method. The principal idea behind this method is to define a measure of distance between the final-state particles and iteratively cluster parti-

cles into pseudojets until some criteria is satisfied.

$$d_{ij} = \min(k_{t,i}^{2p}, k_{t,j}^{2p}) \frac{R_{ij}^2}{R^2} \quad (3.2)$$

$$R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2 \quad (3.3)$$

$$d_{iB} = k_{t,i}^{2p} \quad (3.4)$$

The equation 3.2 defines the measure of the distance between entities (particles, pseudo-jets) i and j as d_{ij} which depends on the transverse momentum ($k_{t,i/j}$) and the angular separation R_{ij} (defined by the rapidity $y_{i/j}$ and azimuthal angle $\phi_{i/j}$ in Eqn. 3.3) within a radius parameter R .

For the anti- k_t algorithm, the parameter p , is set to -1 . Accordingly, the Eqn. 3.2,3.4 are interpreted as Eqn. 3.5.

$$d_{ij} = \frac{1}{\max(k_{t,i}^2, k_{t,j}^2)} \frac{R_{ij}^2}{R^2} \quad \text{and} \quad d_{iB} = \frac{1}{k_{t,i}^2} \quad (3.5)$$

The algorithm starts with looking for the pair of entities with the smallest d_{ij} . The combination of the pair takes place when $d_{ij} < d_{iB}$. If $d_{ij} > d_{iB}$, then i is called a jet and removed from the list of entities. The algorithm then recalculates the d_{ij} for all possible entities, picks the smallest d_{ij} pair and iterates over the steps until no further recombination is possible because all possible $d_{ij} > d_{iB}$. For the anti- k_t algorithm, practically, this makes it more likely that all soft particles j within a distance parameter R would iteratively get added to the cluster around a neighbouring hard particle 1 since the d_{ij} between soft particles would be much larger as compared to d_{1j} (and hence not favoured for the combination check). As a result, the anti- k_t algorithm favours the production of cone-like jets that are built around the particles with the highest p_T and also ensure infrared and collinear safety.

At the CMS experiment, the default angular size for jets (defined by the distance parameter R) is taken to be 0.4 and are referred to as AK4 jets. Several analyses target boosted topologies where there is overlap between the individual objects (for eg, boosted jets) in which case jet algorithms with $R = 0.8$ are often preferred. Depending on the objects defined as the particle candidates for

the jet clustering algorithms, there are several options available. In most cases, jets clustered from particle-flow candidates are used. Due to the contribution of secondary particles coming from pileup, an additional step is included to mitigate the impact of pileup in the reconstructed jets. During the Run 2 data-taking of the CMS experiment, the standard pileup mitigation technique has been the Charged Hadron Subtraction (CHS) algorithm. The CHS algorithm removes the charged particles associated to pileup vertices before the clustering algorithm using the information provided by the tracker. This however suffers from limited coverage as the tracker extends until $|\eta| = 2.5$. An alternative approach to this has been developed called PileUp Per Particle Identification (PUPPI) [131, 132]. The PUPPI algorithm evaluates a probability for each particle (charged or neutral) based on the information from the surrounding particles and assigns higher weights to those originating from the primary vertex. For Run 2 data-taking, jets reconstructed with the CHS algorithm (CHS jets) have been extensively used and are the recommended jet objects. For the Run 3 data-taking, the PUPPI algorithm has been adopted as the preferred method of pileup mitigation.

Jet energy resolution and corrections

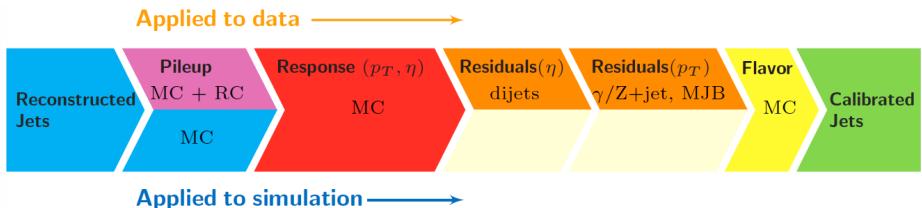


Figure 3.6: Factorized steps of jet energy corrections are shown for data (top row) and for simulation (bottom row). Each step corrects for a specific effect and is applied to jets in CMS. Figure taken from Ref. [133].

Following the jet reconstruction process, the jet energy and momentum still require some additional steps to represent the true four-momenta of the originating parton. There are several contributing factors like presence of pileup and/or non-linear effects of the detector which shift the reconstructed jet energy (termed as raw jet energy) away from the true energy of the parton. These cal-

ibration steps fall broadly under the jet energy scale (JES) and resolution (JER) corrections. The CMS experiment uses a factorized approach for applying the jet energy corrections, as shown in Fig. 3.6. Each step of the correction accounts for a specific effect and are discussed in the following subsections. The final flavor-dependent corrections are optional and are not discussed here. This section focuses on the jet energy resolution and correction results for the CMS Run 2 legacy dataset [133] which is the predecessor to the ultra-legacy datasets.

Pileup offset correction

During the jet reconstruction, energy deposits from pileup interactions can contribute to the reconstructed energy of jets originating from hard-scattering processes. A pileup offset correction is applied on the raw jet energy as the first step of the JES corrections. This correction is derived using reconstructed jets that are matched to particle-level jets in simulation, by comparing their p_T with and without pileup. The average p_T difference between with and without pileup scenarios for jets binned in energy density ρ , jet p_T , η , area, is defined as the pileup offset. The offset is then subtracted from the raw jet energy for all jets in the event. The corrections are determined from QCD multijet simulation samples. While the offset is calculated in simulation, it is applied to both data and simulation. As a result, data requires an additional correction ('residual' correction) which is measured using the random cone method [134]. This method measures the noise and pileup contributions by using events with no hard-scattering processes in data (i.e. zero-bias events). The pileup offset correction step is important for the CHS jets. For PUPPI jets, this correction step is not applied due to the improved pileup suppression provided by the PUPPI algorithm.

Response corrections

Following the pileup correction, the reconstructed jet energy may still vary from the particle-level energy due to detector effects like imperfect energy deposition in the calorimeters. The response correction accounts for such effects and plays a crucial role in the jet calibration procedure. Similar to the pileup offset correction, the corrections in this step are derived from multijet QCD simulations by

comparing the particle-level and reconstructed jet energy for matched jets. The mean jet response, defined as $\langle R_{\text{jet}} \rangle = \frac{\langle p_{\text{T}}^{\text{reco}} \rangle}{\langle p_{\text{T}}^{\text{ptcl}} \rangle}$ where ptcl refers to the particle-level and reco is the reconstructed counterpart, measures the deviation in the jet energy after the pileup offset correction is applied. It is derived in bins of $|\eta^{\text{jet}}|$ and $p_{\text{T}}^{\text{jet}}$ for each era of Run 2 data-taking. The correction is the inverse of the jet response and is propagated to the jet energy as a multiplicative factor. The simulated jet response using the pileup offset corrected jets is shown in Fig. 3.7 (left) for the 2018 simulation. In the barrel region, the simulated response is observed to be stable for different $p_{\text{T}}^{\text{ptcl}}$ ranges for the jets. Towards larger η , the response becomes increasingly dependent on the jet p_{T} . The strongest deviations in response occurs in the EC2 region ($2.5 < |\eta| < 3$) of the endcap calorimeters since the calorimeter coverage enters a transition region to the hadronic forward calorimeter (HF) and is also affected by calorimeter aging. The response after the corrections have been applied are shown in Fig. 3.7 (right) for the 2018 simulation. It is seen that the corrections so far bring the agreement between $p_{\text{T}}^{\text{ptcl}}$ and $p_{\text{T}}^{\text{reco}}$ to $\sim 1\%$ for all η/p_{T} ranges considered, with 0.1% agreement in the barrel and EC1 region of the endcaps.

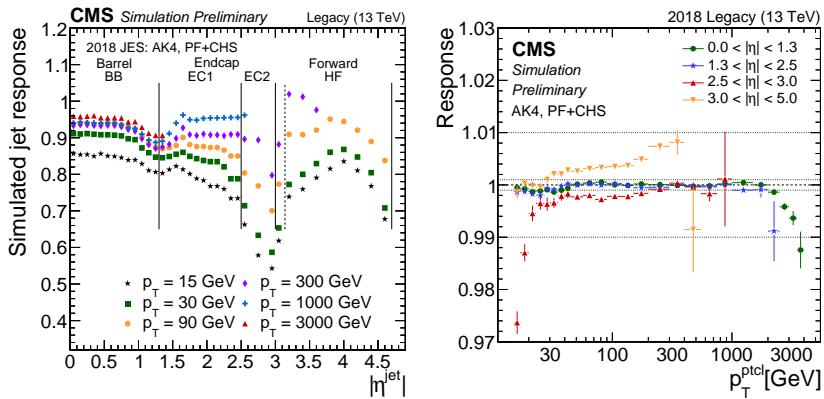


Figure 3.7: The jet energy response before application of the scale corrections (left) and after applying the corrections (right) for various jet η and p_{T} . The corrections lead to $\leq 1\%$ agreement between particle-level and reconstructed jet energy in simulation. Figure taken from Ref. [133].

Residual corrections

Following the response corrections, the jets in simulation reach a good level of closure with the particle-level jet energy. However, there are still some residual differences that remain between data and simulation, indicating that further corrections to the jet energy scale are required for data. These residual corrections are applied in two stages: an $|\eta|$ -dependent correction to correct for different response in the various sub-detectors; followed by a p_T -dependent correction to account for the absolute scale difference in the barrel region. The corrections are obtained using in-situ methods, which rely on well-calibrated reference objects within the event to calibrate the jets. Since data has the previous corrections based on the simulation, the residual corrections are an additional correction on top of it and are applied as a simulation-to-data (MC/data) scale factor. There are two methods based on the above-mentioned idea which are used to derive the corrections - the direct balance (DB) or p_T balance method and the missing transverse momentum projection fraction (MPF) method [134]. The details of each method are discussed in the next Chapter 4.

For calculating the $|\eta|$ -dependent corrections, the MPF method is used in events containing two approximately back-to-back jets in the hard-scattering process (i.e. dijet events in QCD). Fig. 3.8 (left) shows the $|\eta|$ -dependent correction for 2018 data. While the correction is small for the barrel region, the EC2 region of the endcap (transition region) and the HF region show larger deviations. The absolute p_T -dependent corrections consider different hard-processes, namely, Z/γ +jet, QCD multijet and hadronic W decays from $t\bar{t}$ processes. Both DB and MPF methods are implemented, and the corrections are shown in Fig. 3.8 (right).

Jet energy resolution

The final step of the jet calibration procedure consists of accounting for the jet energy resolution (JER) differences between simulation and data. While the JES is the mean of the response distribution, the JER can be quantified as the spread of the response distribution. After the jet energy corrections have been applied, this corresponds to the width of the distribution centered around unity. The de-

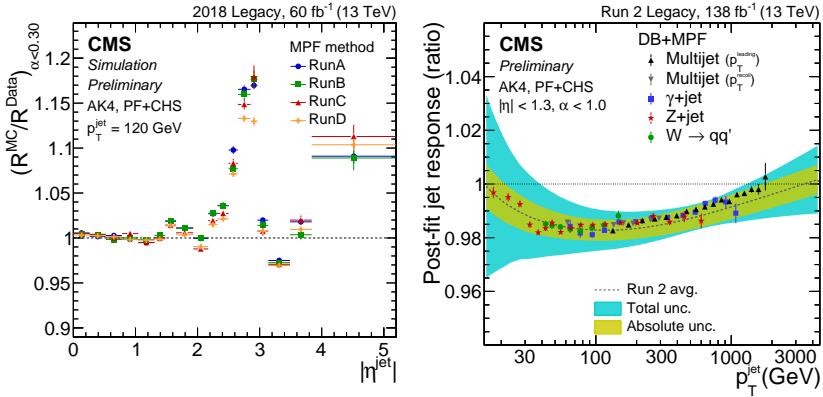


Figure 3.8: The jet energy response before application of the scale corrections (left) and after applying the corrections (right) for various jet η and p_T . The corrections lead to $\leq 1\%$ agreement between particle-level and reconstructed jet energy in simulation. Figure taken from Ref. [133].

tails of the jet energy resolution in Run 2 and related studies are discussed in detail in the next Chapter.

Jet tagging

In addition to reconstructing and calibrating the jets, dedicated algorithms have been developed to allow for the identification of the parton origin in jets. This is called jet tagging and is an essential component in several physics analysis. One such example is top physics where b-tagging (i.e. the selection of jet originating from a b-quark) is required to select the dominant $t \rightarrow b + W$ decay mode for the top quark. Since the b and c hadrons can travel a measurable distance before decaying, the presence of displaced particles within a jet is one of the discriminating variables which allows for b-tagging and c-tagging at the CMS experiment.

The CMS experiment employs several sophisticated algorithms for jet tagging, many of which are based on state-of-the-art machine learning models. The complex internal structure of jets and the subtle differences between heavy flavour and light flavour jets make them ideal for deep learning approaches, which excel at pattern recognition problems. During Run 2, the DeepCSV and

DeepJet algorithms [135, 136] have been used to provide a high heavy-flavour tagging efficiency while maintaining a low mistagging rate of light-flavour jets.

3.2.8 Missing transverse momentum

The head on nature of the pp collisions implies that there is no net transverse component initially. Accordingly, the net transverse momentum in the event is expected to be zero. However, this is the case when all outgoing particles are reconstructed. In processes involving neutrinos, it is expected that there will be some transverse imbalance present since they pass through the detector without any interaction. Similarly, in several BSM theories, weakly-interacting particles are often proposed which would lead to events with ‘missing’ energy. Therefore, the missing transverse momentum forms an important object for various physics analyses. It is defined as the negative vector sum of the p_T of all reconstructed PF objects (Eqn. 3.6).

$$\vec{p}_T^{\text{miss}} = - \sum^{\text{PF objects}} \vec{p}_T \quad (3.6)$$

Similar to other physics objects, the p_T^{miss} requires some dedicated corrections. One of the dominant sources of inaccuracy in p_T^{miss} arises from the nonlinear response of calorimeters. Since this effect is already corrected for jets with the calibration steps outlined before, the p_T^{miss} measurement is improved by propagating the p_T correction of the jets to the raw p_T^{miss} . This procedure also helps in reducing the effect of pileup and results in a p_T^{miss} response close to unity [137].

Chapter 4

Jet energy resolution in the CMS experiment

Jets reconstructed in the CMS experiment undergo a multistep energy calibration procedure (Fig. 3.6) to correct for various effects such as pileup contamination and non-linear detector response, in both simulation and data. A central quantity used to evaluate the accuracy of the jet reconstruction is the jet response, defined as:

$$R_{\text{jet}} = \frac{p_T^{\text{reco}}}{p_T^{\text{true}}} \quad (4.1)$$

In the case of the simulation-based response corrections discussed in Section 3.2.7, $p_T^{\text{true}} = p_T^{\text{ptcl}}$ where p_T^{ptcl} is the transverse momentum of the particle-level jet that is matched to the reconstructed jet. For the jet energy scale, the mean of the response, $\langle R_{\text{jet}} \rangle$, is used to estimate the correction. The jet energy resolution, on the other hand, is described as the spread/width of the response as mentioned before in Section 3.2.7.

To get a first glance at the jet energy resolution (JER) in Run 2 datasets, it has previously been studied using QCD multijet simulations [133] by evaluating the width of the jet response distribution in bins of jet p_T , η and number of pileup interactions (μ). Fig. 4.1 shows the simulated JER as a function of jet p_T in both the central and forward η regions, across different pileup conditions. The resolution in the barrel region for $p_T > 100$ GeV is stable and roughly 5–10%.

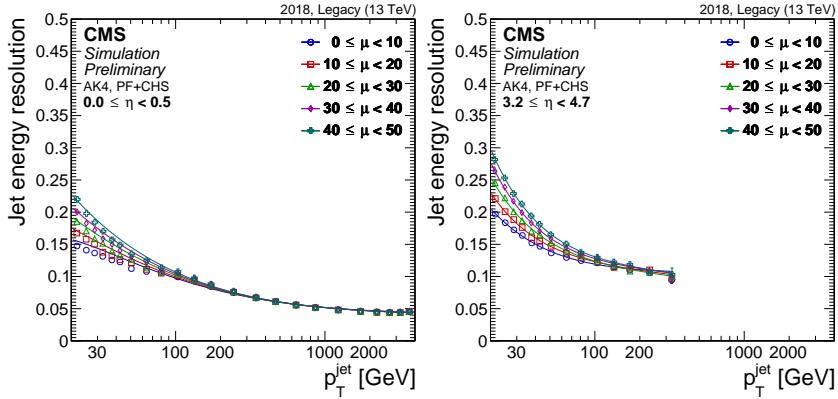


Figure 4.1: The jet energy resolution in simulation in the barrel region for $0 \leq |\eta| \leq 0.5$ (left) and in the HF region for $3.2 \leq |\eta| \leq 4.7$ (right). Figure taken from Ref. [133].

It is also seen that as the jet p_T decreases, the jet energy resolution worsens as it becomes harder to distinguish pileup contribution in the jet originating from the purely hard-process. As a result, for $p_T < 100$ GeV, the JER becomes pileup dependent. This is also seen in Fig. 4.2 where the jet energy scale uncertainty is shown as a function of the jet p_T for a central jet ($|\eta^{\text{jet}}| = 0$) in the 2018 dataset. Among the several factors contributing to the total scale uncertainty, pileup becomes a significant source of uncertainty when jet p_T drops below 50–100 GeV.

For comparisons between data and simulation, in-situ methods, as briefly mentioned for the residual corrections (Section 3.2.7), can be used to derive the jet energy resolution for both data and simulation. Fig. 4.3 shows the JER data/MC scale factors as a function of the $|\eta|$ for jets with $p_T > 100$ GeV. The JER scale-factors are found to be within 1 to 1.2 except for the endcap to HF transition region where larger deviations are observed. The p_T -dependent results of the scale-factors are shown in Fig. 4.3 for $0.522 \leq |\eta| \leq 0.783$ (barrel region) and $2.853 \leq |\eta| \leq 2.964$ (endcap to HF transition region) with Z+jets and QCD dijet events. Besides $2.5 \leq |\eta| \leq 3$, it is found that scale-factors are constant with respect to the jet p_T (~ 1.1 for barrel region). This leads to an important observation that the resolution in data is typically worse than what is found in

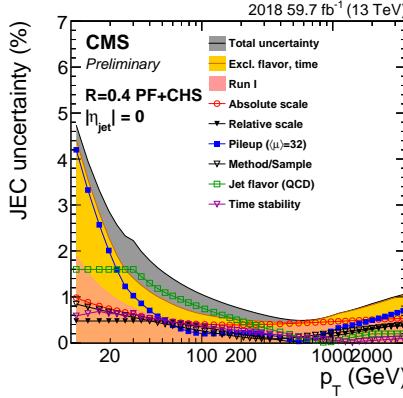


Figure 4.2: Dominant p_T -dependent jet energy correction uncertainties for 2018 data set. Pileup uncertainty contributes significantly below 50 GeV. Figure taken from Ref. [133].

simulation. An accurate modeling of this effect in simulation is essential, as it plays a critical role in analyses involving steeply falling spectra or resonance decays, and impacts the description of event observables such as missing transverse energy (MET). The broader resolution in data is propagated to simulation using jet smearing techniques such as the stochastic and scaling methods, which is discussed later in this section.

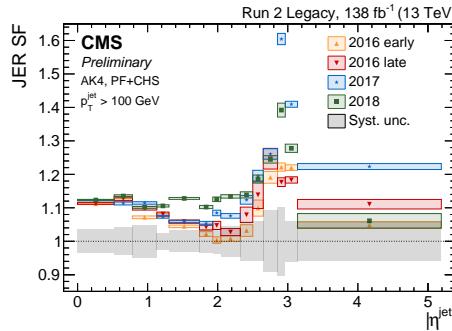


Figure 4.3: The η -dependent jet energy resolution scale-factors for Run 2 datasets for jet $p_T > 100$ GeV. The scale-factors are within 1–1.2 for all regions except for the endcap-HF transition region. Figure taken from Ref. [133].

The jet energy results discussed so far have been derived for the Run 2 data-

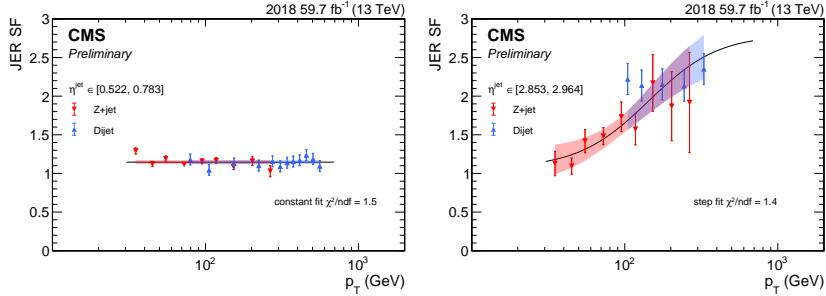


Figure 4.4: The p_T -dependent jet energy resolution scale-factors in the barrel region for $0.522 \leq |\eta| \leq 0.783$ (left) and in the HF region for $2.853 \leq |\eta| \leq 2.964$ (right) for the 2018 legacy datasets. Figure taken from Ref. [133].

taking periods, namely in years 2016, 2017 and 2018. The collected datasets have been passed through several processing campaigns to improve the quality of the simulations and data with the latest available detector calibrations, algorithms and corrections. The results presented here so far have been derived for the legacy processing of the data, which is the penultimate processing campaign for Run 2. The final and the most refined reprocessing for Run 2 is the ultra-legacy campaign. Consequently, with the improved data quality from this campaign, it is vital to quantify the jet performance using this dataset. The present Chapter aims to assess the jet energy resolution in the ultra-legacy datasets of Run 2 with an emphasis on jets with $p_T < 100$ GeV, and is organized as follows. Sections 4.1–4.4 present a new method for estimating the jet energy resolution in data and simulation, employing template-based fits to model pileup and additional contributions. The results of the study for Run 2 are shown in Section 4.5. A complementary, recently developed method based on missing transverse energy [138] within the CMS Collaboration is applied, and its results are presented in Section 4.6. Section 4.7 presents some validation studies to compare the performance of both methods.

4.1 In-situ method for jet response: p_T -balance

The idea behind the in-situ methods is to select hard-scattering processes consisting of a reference object and the uncorrected jet (typically recoiling against

each other) and using the transverse momentum balance in the event to derive the jet response at a reconstructed level. This enables measuring the correction and, in this case, the resolution for simulation and data separately. As a result, a data/MC scale factor is obtained and used to smear the jets in simulation. The approach adopted here is called the p_T -balance method.

The p_T -balance approach evaluates the jet response by considering a better-calibrated reference object like a reconstructed Z and directly comparing to the reconstructed jet p_T . The p_T balance response, $R_{\text{jet}}^{\text{balance}}$, is defined as shown in Eqn. 4.2

$$R_{\text{jet}}^{\text{balance}} = \frac{p_T^{\text{reco}}}{p_T^{\text{ref.}}}, \quad \text{ref.} \in [\text{Z}, \gamma, \text{jet}] \quad (4.2)$$

where p_T^{reco} and $p_T^{\text{ref.}}$ are the reconstructed p_T of the jet and the reference object, respectively. There are several candidates that can be considered as the reference object. The reconstructed Z is an ideal reference candidate, in particular for the $Z \rightarrow \mu\mu$ decay channel, as the muon scale is calibrated up to $0.1 - 0.3\%$ [126]. Similarly, the reconstructed photon can also be used as a reference object as in the case for the residual corrections discussed before. The most popular choice for the jet energy resolution studies is the dijet event topology where two jets recoiling against each other are selected to measure the resolution from the asymmetry distribution $\left(A = \frac{p_T^{\text{jet 1}} - p_T^{\text{jet 2}}}{2p_T^{\text{avg}}} \right)$. To ensure at least one well-calibrated jet in the dijet event, one jet is required to lie in the barrel region ($|\eta| < 1.3$) because of the small variation in the response there [139]. It is important to note that each of the objects provide sensitivity to different regions of the jet p_T spectrum. The Z+jet events provide access to jet p_T up to 200 GeV as seen in Fig. 4.4. The dijet event topology takes over for jet $p_T > 100$ GeV with complementarity provided by γ +jet events which are sensitive to a similar p_T range. Combining the different channels are therefore vital to get a complete description of the jet energy response across a wide range of jet p_T . A schematic illustrating a typical Z+jet event is shown in Fig. 4.5.

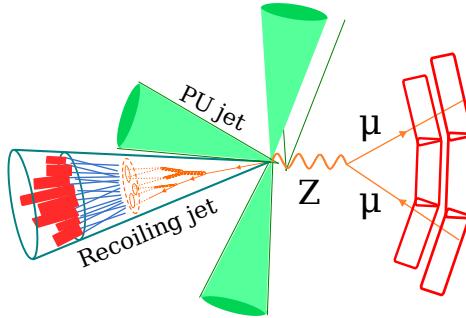


Figure 4.5: Schematic illustrating the p_T -balance topology with $Z+jet$ events.

4.2 Data and simulated samples

The data used in this analysis is from proton-proton collisions recorded at $\sqrt{s}=13$ TeV in 2018 corresponding to 59.8 fb^{-1} . All data and MC samples used in the analysis are ultra-legacy (UL) Run 2 samples which is the final reprocessing of the Run 2 datasets and incorporates all known improvements in detector alignment, reconstruction etc. The analysis targets $Z+jet$ events in data and simulation, in the $Z \rightarrow \mu\mu$ final-states. The following trigger is considered: a dimuon trigger that requires two isolated muons with $p_T > 17$ and 8 GeV respectively, forming $m_{\mu\mu} > 3.8 \text{ GeV}$. In addition to this, a condition on the longitudinal impact parameter with respect to the leading vertex is applied ($dz < 0.2 \text{ cm}$).

Various MC samples for the Drell-Yan process simulation have been used. The nominal sample consists of Z bosons produced in association with ≤ 2 jets at leading order (LO) accuracy in the QCD coupling, using the event generator MadGraph5 [97, 98]. This is the primary sample used for the jet energy resolution measurements discussed in this Chapter. For validation studies, two next-to-leading order (NLO) samples are considered which have been generated using MadGraph5_aMC@NLO [99, 100] and POWHEG [94–96]. While the LO sample is a less-accurate simulation compared to others that have been considered, it provides a large sample size and is sufficient to model the event topology (i.e. a jet (typically the leading p_T jet) recoiling against the Z). These samples, generated at the matrix-element level, are followed by parton shower and hadronisation processes modelled with PYTHIA8.240 [103], using the

CP5 underlying event tune [140]. Additionally, for the MadGraph5 and MadGraph5_aMC@NLO samples, events with dilepton mass greater than 50 GeV are generated. For the POWHEG sample, events with dilepton mass within $50 - 120$ GeV have been generated. Since a rich-purity of Z+jet events can be obtained from data through tight selection cuts, no other background sample has been considered. The complete list of MC samples are listed in Table 4.1.

Process	ME Generator	σ (pb)
DY+jets ($m_{ll} > 50$ GeV)	MadGraph5 (LO)	5398
DY+jets ($m_{ll} > 50$ GeV)	MadGraph5_aMC@NLO (NLO)	6424
DY+jets ($120 > m_{ll} > 50$ GeV)	POWHEG (NLO)	2116

Table 4.1: Simulation samples used in the analysis.

4.3 Event selection

The Z+jet event topology consists of jets and muons/electrons, reconstructed by the particle flow algorithm [107]. The jets are reconstructed using the anti- k_T algorithm with a radius parameter of $R=0.4$. To reduce the impact of pileup in the jet reconstruction, pileup mitigation algorithms such as Charged Hadron Subtraction (CHS) or PileUp Per Particle Identification (PUPPI) are additionally applied. For this analysis, only jets passing the CHS algorithm have been considered. The jets are then calibrated using the latest jet energy corrections according to the factorized approach as discussed in Section 3.2.7. As the missing E_T (i.e. p_T^{miss}) is the negative sum of all activity in the event, it is sensitive to changes in the jet energy after calibration. Accordingly, the p_T^{miss} is modified to include the jet energy after the calibration instead of the uncorrected jet energy. While no requirements have been placed on the p_T^{miss} in the event, this step is required for the complementary method to the p_T balance approach, discussed later in Section 4.6. The jets are also required to pass the jet identification criteria, based on jet properties such as the neutral and charged hadron energy fractions, neutral and charged multiplicity etc, to ensure a high-quality jet. In addition to this, any jet that overlaps with a muon or photon (determined by ΔR between jet and nearest muon or photon) is removed to reduce fake jets. All jets with $p_T > 12$ GeV

are stored in the jet collection.

For the Z boson candidate, the $Z \rightarrow \mu\mu$ channel is considered as the primary decay mode as it provides the largest statistics due to the lower p_T thresholds for the muons and a better Z mass resolution. The muons pass a set of quality cuts to select well-reconstructed Z bosons. The dimuon trigger is used to select the events as mentioned in the previous section. In the muon channel, events containing two muons passing the tight identification [126] and tight isolation criteria are stored. The leading and the subleading muons are additionally required to have p_T greater than 20 and 15 GeV respectively. The reconstructed Z boson candidate is selected within the dimuon mass of 70 to 110 GeV. No explicit p_T cut is applied on the reconstructed Z boson, ensuring that all p_T^Z ranges remain accessible. One such use case is modeling the p_T balance response distribution for pileup jets (discussed later). However, the JER measurement itself will be carried out in bins of p_T^Z , with the first bin starting from 30 GeV.

The ideal Z+jet event topology consists of the hard-scattering jet recoiling against the Z boson, ensuring a good balance and an accurate jet response description. This is most often not the case due to the presence of additional jet activity arising from underlying event but also from pileup. A common approach for measuring the response in the ‘ideal’ topology is termed as the α extrapolation method. The additional hadronic activity in the event is quantified by $\alpha = \frac{p_{T,2nd\ jet}}{p_{T,ref.}}$ where the p_T of the second highest p_T jet in the event is compared to the p_T of the reference object (in this case, the Z boson). A large α indicates significant additional soft activity in the event, coming from pileup or radiation effects. Such events can distort the jet response distribution since the leading jet may not perfectly recoil against the Z boson. By applying a maximum cut on the α (α_{max}), events with large additional activity are rejected. The response distribution is then studied to extract the jet energy resolution. Furthermore, the response is measured at different values of α_{max} to linearly extrapolate to the $\alpha \rightarrow 0$ scenario (the ideal topology) [134].

There are two major challenges that arise with this method when moving towards jets of $p_T < 50$ GeV. The first obstacle arises from the increasing influence of pileup on the response as seen in Fig. 4.2. This is because it becomes increasingly difficult to differentiate between the hard-process jet and pileup jets,

as the p_T range begins to overlap. Furthermore, the α extrapolation method is inherently biased towards low pileup. This can be illustrated by checking the distribution of the number of primary vertices with the application of a particular α cut. Fig. 4.6 (left) shows this distribution (red triangles) for $p_T^Z \in [30, 50]$ GeV where events with only one jet with $p_T > 10$ GeV are allowed. This approximately corresponds to $\alpha = 0.3$ which is one of the working points used to perform the extrapolation [133]. This is compared to the inclusive case where events are allowed to have one or more jets with $p_T > 10$ GeV (black dots) for any p_T^Z . In this analysis, a loose set of selection criteria is applied based on the jet activity in different kinematic and geometric regions as described below. Events satisfying the following selection criteria are selected:

- At most 1 low p_T central jet in $|\eta^{\text{jet}}| < 2.5$, ≤ 1 jet with $p_T \in [12, 20]$ GeV passing the pileup jet identification criteria [141] is allowed: As illustrated in Fig. 4.6 (left), additional jets can arise due to higher pileup with $p_T > 10$ GeV. The pileup jet identification helps in suppressing these events.
- At most 1 central jet with $p_T > 20$ GeV in $|\eta^{\text{jet}}| < 2.4$, ≤ 1 jet with $p_T > 20$ GeV is allowed: Most pileup jets in the central region are expected to have $p_T < 20$ GeV.
- At most 1 jet with $p_T > 30$ GeV in $|\eta^{\text{jet}}| < 5.191$, ≤ 1 jet with $p_T > 30$ GeV is allowed: Beyond the central region, there are a larger number of CHS jets coming from pileup. A higher threshold for candidate jet is accordingly applied.

The selection criteria presented above lead to low additional jet activity in the events to ensure a good balance between the probe jet and the Z boson. It also reduces the bias due to pileup as can be seen in Fig. 4.6 (left) where the distribution from events passing the selection criteria (blue squares) is closer to the distribution where additional jets with $p_T > 10$ GeV are allowed.

An additional challenge in the low- p_T regime is related to the hadronic activity arising from underlying event and multi-parton interactions. One of the primary contributions come from initial-state radiation (ISR) and final-state radiation (FSR) which could lead to a worse response distribution between the Z

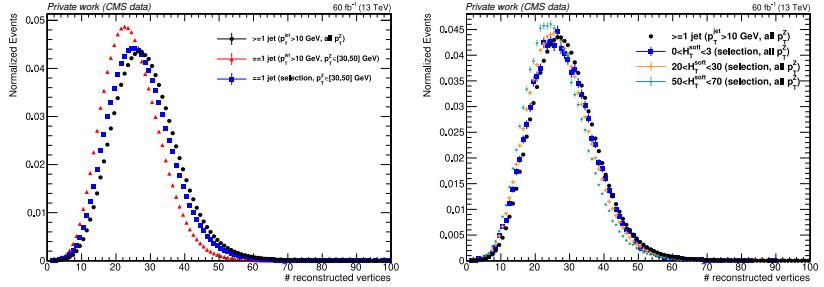


Figure 4.6: (Left) Distribution of the number of reconstructed vertices for $Z + \text{jet}$ events with ≥ 1 jets, $== 1$ jet and $== 1$ jet after passing the selection to illustrate that events passing the selection have reduced bias towards less pileup. (Right) Distribution of the number of reconstructed vertices for events in different H_T^{soft} ranges to illustrate that H_T^{soft} introduces negligible bias towards lower pileup.

boson and the jet. Furthermore, ISR and FSR are not accurately modeled in simulations and their description varies among different simulation frameworks. This provides the motivation to add a selection criterion based on the amount of ‘soft’ hadronic activity in the event, denoted by H_T^{soft} . The particles arising from the underlying event are typically very soft (i.e. low p_T) which may lead to unclustered energy deposits in the event. However, the charged hadrons arising from this appear as tracks in the event since the CMS tracker reconstructs charged particles with p_T as low as 0.3 GeV [108]. The hadronic activity in the event can be measured by calculating the scalar sum p_T of tracks in the event coming from the primary vertex (defined as

$H_T = \sum_{\pi^\pm \text{ tracks}} p_T$ (tracks from PV)). The H_T^{soft} is then evaluated by subtracting the p_T contribution from the charged hadronic tracks associated to the probe jet. There is no pileup dependence of H_T^{soft} as can be seen in Fig. 4.6 (right) where the distribution of number of reconstructed vertices is independent for different ranges of H_T^{soft} . The following selection is included in the selection:

- For $p_T^Z < 105$ GeV, $H_T^{\text{soft}} < 15$ GeV
- For $p_T^Z > 105$ GeV, $H_T^{\text{soft}} < 0.5 \times p_T^Z$: Larger H_T^{soft} threshold is needed since the ISR and FSR contribution increases with probe jet p_T .

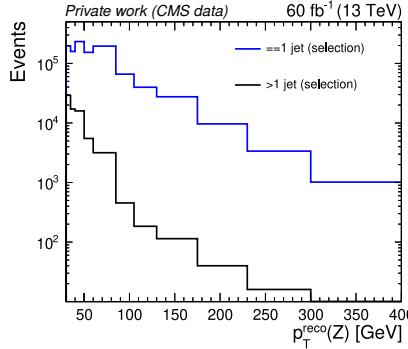


Figure 4.7: Reconstructed Z p_T distribution events with only one jet (blue) and multiple jets (black) after the selection is applied.

Following the full event selection, there are still multiple jets left in each event. In the jet collection mentioned before, an additional criterion is included to exploit the back-to-back nature of the ideal Z+jet topology. Jets with $\Delta\phi(Z, \text{jet}) > 3$ are stored as probe jets. Fig. 4.7 shows the p_T^Z for the events with one probe jet only (blue) and with multiple probe jets (black). It is observed that a substantial portion of the dataset has multiple probe jets in the event and this contribution increases as the p_T^Z decreases.

4.3.1 JER extraction

The response of the detector is initially measured in bins of p_T^Z and $|\eta|$ of the jet. Fig. 4.8 shows the p_T balance response in simulation in the central (left) and the forward (right) region in the low p_T^Z bin. Since the jets undergo the scale correction, the response distribution is centred around 1 in both cases. The response distribution can be split into two components – the jets from the hard-scattering process (i.e. gen. matched jets) and pileup (i.e. non-gen. matched jets). A reconstructed jet is defined as a gen. matched jet when a generator-level jet is present within a cone of $\Delta R = 0.2$. It is seen that there is a substantial contribution coming from pileup (red), especially for the forward region, as indicated by the non-gen. matched jets (green). Eqn. 4.3 shows the response distribution as a function of these two components, defined as the pileup and the recoiling jet.

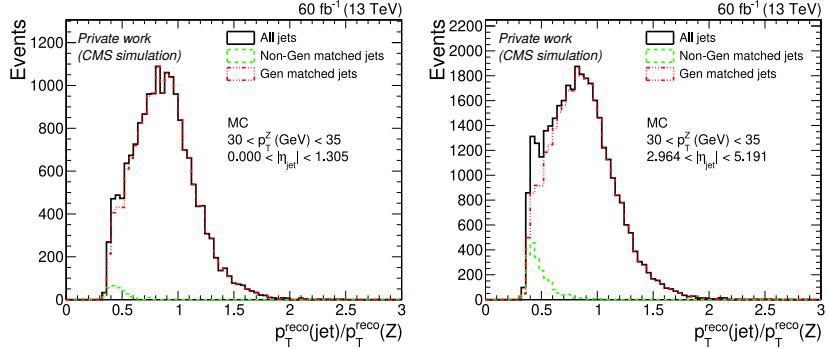


Figure 4.8: p_T balance distribution in simulation for $p_T^Z \in [30, 35]$ GeV in the central region (left) and the forward region (right). The gen. matched jets are shown in red and the non-gen. matched jets (i.e. pileup jets) are shown in blue.

Each of these contributions require dedicated modelling as discussed below.

$$\frac{p_T^{\text{reco}}(\text{jet})}{p_T^{\text{reco}}(Z)} = \frac{p_T^{\text{reco}}(\text{pileup jet})}{p_T^{\text{reco}}(Z)} \oplus \frac{p_T^{\text{reco}}(\text{recoiling jet})}{p_T^{\text{reco}}(Z)} \quad (4.3)$$

Pileup template

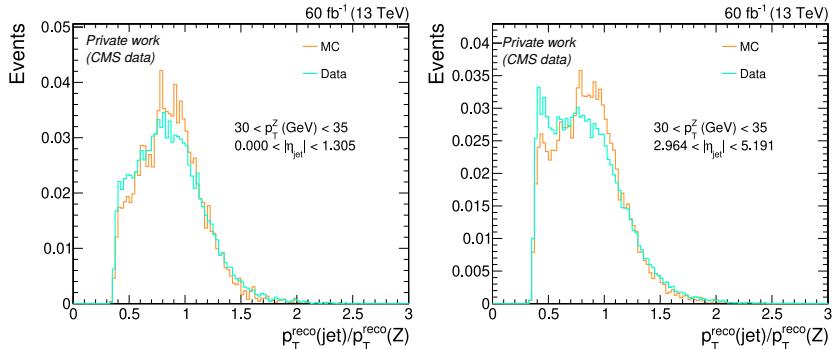


Figure 4.9: p_T balance distribution in data and simulation for $p_T^Z \in [30, 35]$ GeV in the central region (left) and the forward region (right). Larger contribution for data is seen at small p_T balance response values compared to simulation.

One of the straightforward approaches to modelling the pileup is to use the non-gen. matched jets in simulation to model the pileup in data. However, this

assumes that the same pileup contribution is expected in data. This is not the case as shown in Fig. 4.9 in the low- p_T region ($p_T^Z \in [30, 35]$ GeV). There is a larger proportion of events at the lower end of the response in data as compared to the simulation. This further provides the motivation to take an ‘in-situ’ approach to estimate the pileup. This approach is based on constructing a template using pileup jets identified in Z +jet events passing similar selection criteria as before.

The first stage of the method involves selecting pileup jets from low energy Z +jet events. The same selection criteria as mentioned in the previous section is used to reconstruct the Z boson candidate except a tighter mass window is applied ($m_Z^{\text{reco}} \in [80, 100]$ GeV). The events are also required to pass the same selection criteria on the number of jets and the hadronic activity in the event as before. All events with $p_T^Z < 5$ GeV are selected to collect events rich in low p_T jets and pileup jets. The jets are required to pass the same identification and p_T criteria as mentioned before. To build the pileup jet collection, only the jets near the Z boson are collected by requiring $\Delta\phi(Z, \text{jet}) \in [1, 2]$ condition. This is in contrast to the probe jet collection mentioned before, which consists of jets with $\Delta\phi(Z, \text{jet}) > 3$. This selection leads to a group of pileup jets which is then used to artificially build Z +jet events in the next stage.

For all events passing the event selection criteria to measure the response distribution, a corresponding Z +jet event for pileup is constructed by using the reconstructed Z boson in the event and associating a random jet from the pileup jet collection generated in the previous stage. This artificial Z +jet event is then used to build the p_T balance response distribution for pileup. To increase the statistics, this is performed multiple times for the same Z boson. This method is applied to both simulation and data to construct the pileup templates in bins of p_T^Z and $|\eta^{\text{jet}}|$.

Figure 4.10 shows the pileup templates for data, simulation and the response distribution for the non-gen. matched jets in a sample bin. It is observed that there is a good agreement between the non-gen. matched jets and the pileup template generated for simulation. This validates the method described here and also increases the statistics available for modelling the pileup. There are also small differences between pileup templates for data and simulation. Therefore, the pileup template is generated separately for simulation and data, and applied

to each sample accordingly.

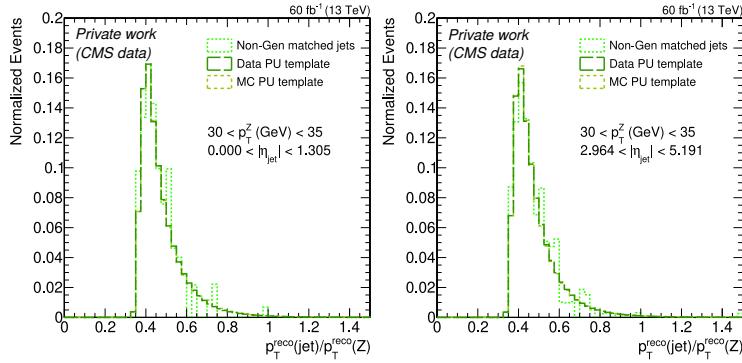


Figure 4.10: Comparison of pileup templates derived from simulation (small dashed histogram) and from data (large dashed histogram) with the distribution of non generator matched jets in simulation (dotted histogram).

Response for recoiling jets

The second term in Eqn. 4.3 contains the contribution from the jets recoiling against the Z boson. As stated in Ref. [134], the response for the recoiling jet can be further divided into three contributing terms. This is shown in Eqn. 4.4.

$$\frac{p_T^{\text{reco}}(\text{recoiling jet})}{p_T^{\text{reco}}(Z)} = \frac{p_T^{\text{gen}}(\text{jet})}{p_T^{\text{gen}}(Z)} \frac{p_T^{\text{gen}}(Z)}{p_T^{\text{reco}}(Z)} \frac{p_T^{\text{reco}}(\text{jet})}{p_T^{\text{gen}}(\text{jet})} \quad (4.4)$$

The first term relates to the response at the generator-level between the Z boson and the recoiling jet. Accordingly, it is termed as the particle-level imbalance (PLI) and originates from the underlying event, initial and final state radiation processes. This is shown in Fig. 4.11 (left) for a representative p_T^Z bin. A template for the PLI is taken directly from simulation and applied when modelling the p_T balance response distribution in data and simulation. This assumes that the PLI is the same in both. A systematic study of the effect due to this choice is studied later in Sec. 4.4. The second term in Eqn. 4.4 is the response for the Z boson momentum. The muons used to reconstruct the Z boson are well-calibrated objects, with a p_T resolution of 1 – 3% [126]. This leads to a very narrow resolution for the detector response of the Z boson that is also centered

at unity (scale-corrected), as seen in Fig. 4.11 (middle). It is hence assumed that the contribution to the recoiling jet distribution resolution for this is negligible and has not been considered.

The third term in Eqn. 4.4 denotes the detector response of the jet. Since the detector response generally arises from calorimeter noise, this can be approximately described by a Gaussian distribution. This is also evident from Fig. 4.11 (right), where the detector response is shown for a representative bin.

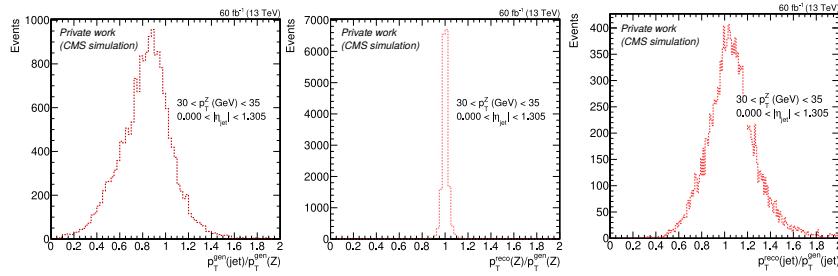


Figure 4.11: Different components used to model the response of the recoiling jets in simulation: particle-level imbalance (left), Z-boson p_T resolution (middle), and true jet response distribution (right).

4.3.2 Scale factor derivation

Combining the different components of the jet response and the pileup contribution, the p_T balance response distribution is modelled by Eqn. 4.5 to estimate the jet energy resolution (σ_{JER}).

$$\frac{p_T^{\text{reco}}(\text{jet})}{p_T^{\text{reco}}(Z)} = \frac{p_T^{\text{reco}}(\text{pileup jet})}{p_T^{\text{reco}}(Z)} \oplus \frac{p_T^{\text{gen}}(\text{jet})}{p_T^{\text{gen}}(Z)} \otimes \mathcal{N}(\mu, \sigma_{\text{JER}}) \quad (4.5)$$

The pileup template generated from the in-situ method, as discussed in Section 4.3.1, is used to derive the binned template-based probability density function (PDF) for the fitting function. The recoiling jet contribution is modelled with the PLI template and a Gaussian with freely-floating mean (μ) and width (σ). The PLI template is converted into a PDF and then convoluted with a Gaussian function to describe the smearing effect of the detector response. For each component of Eqn. 4.5, the normalization is allowed to float. The fit is performed

within a custom fitting range which has a lower bound of $12 / \min(p_T^Z)$ in the bin and a higher bound that excludes 2.5% of the distribution. For bins with $p_T^Z > 50$ GeV, the lower bound is shifted such that 2.5% of the distribution is excluded from the lower end. This is done to remove any effects arising from the tails of the distribution and to ensure a good fit. Additionally, for simulation, half of the dataset is used as the test sample while the other half is used to generate the PLI template.

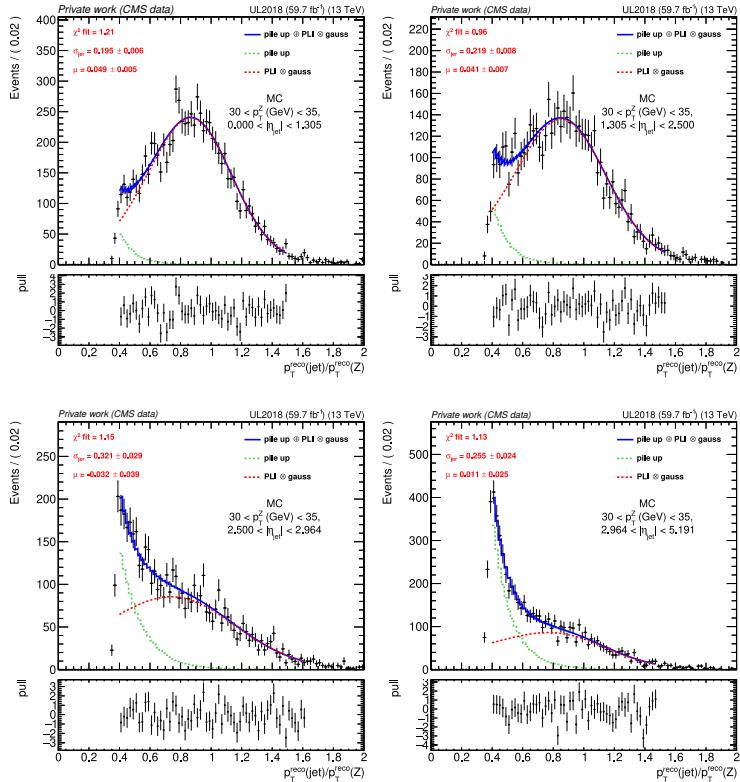


Figure 4.12: p_T balance distribution for $30 < p_T^Z < 35$ GeV in the Barrel (top left), Endcap 1 (top right), Endcap 2 (bottom left) and Forward (bottom right) regions. The JER values are extracted from the fit using Eqn. 4.5.

The jet energy response from the p_T balance response distribution is extracted in bins of p_T^Z and $|\eta^{\text{jet}}|$. The p_T^Z bins are defined by the edges [30, 35, 40, 50, 60, 85, 105, 130, 175, 230, 300, 400] GeV. While a fine $|\eta^{\text{jet}}|$ binning is

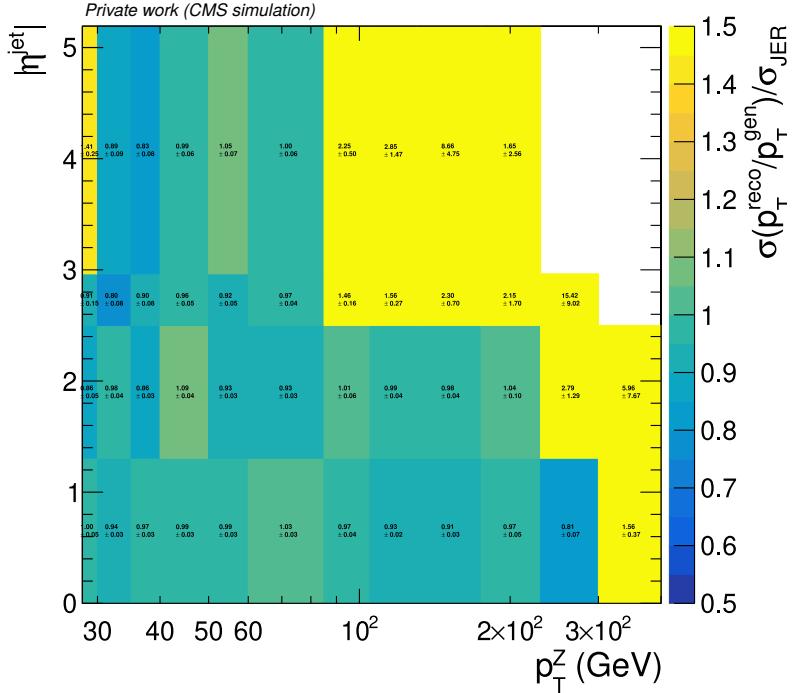


Figure 4.13: The ratio of the true jet energy resolution and the extracted resolution as a function of $|\eta^{\text{jet}}|$ and p_T^Z . A ratio close to unity indicates good closure, demonstrating that the JER extraction method accurately reproduces the true resolution.

used to accurately measure the JER in different detector regions, a coarse binning is used as a first estimate of the JER in the Barrel region (BB: $0 < |\eta^{\text{jet}}| < 1.305$), Endcap 1 region (EC1: $1.305 < |\eta^{\text{jet}}| < 2.5$), Endcap 2 region (EC2: $2.5 < |\eta^{\text{jet}}| < 2.964$) and Forward region (HF: $2.964 < |\eta^{\text{jet}}| < 5.191$). Figure 4.12 shows the p_T balance distribution in simulation and the resulting fit for $30 < p_T^Z < 35$ GeV with the coarse binning. The pileup contribution increases with the $|\eta^{\text{jet}}|$ and the generated pileup template accordingly scales to ensure a stable fit across all the regions. The reduced precision from the tracker and calorimeter at larger $|\eta^{\text{jet}}|$, along with increased sensitivity to soft activity, impact the jet reconstruction performance and lead to larger JER as $|\eta^{\text{jet}}|$ increases.

To study the closure of these results, the resolution extracted from the fit is

compared to the true jet resolution from the simulation (i.e. $\sigma(\frac{p_T^{\text{reco}}(\text{jet})}{p_T^{\text{gen}}(\text{jet})})$). Figure 4.13 shows the ratio in coarse bins of p_T^Z and $|\eta^{\text{jet}}|$. Agreement within 5-10% is observed between the true jet resolution and the extracted resolution from the fit for the Barrel and Endcap 1 regions. Similar agreement is observed for the Endcap 2 and Forward regions after taking into account the larger uncertainties.

Some example p_T balance distributions for data and simulation are shown in Fig. 4.14 for the Barrel region. While there is a slightly larger pileup contribution observed in data, the respective data and simulation pileup templates accordingly ensure stable fits. It is also evident from the relatively simpler no pileup scenario ($105 < p_T^Z < 130$ GeV) that the resolution is larger in data and is reflected in the extracted fits across all bins. The jet energy resolution scale factors in coarse and fine bins are shown in Fig. 4.15. In general, the jet energy resolution for data is observed to be 10-30% larger than the resolution for simulation, with a possible p_T dependence. Accordingly, a p_T -dependent study of the JER scale factors is carried out and described in Section 4.5.

4.4 Systematic uncertainties

The jet energy resolution measurement has several sources of systematic uncertainty, which need to be accounted for in both data and simulation. For each type of uncertainty, the impact on the resolution is measured with the up and down variations. The resulting shift in JER and the scale-factors is then applied as a symmetric uncertainty relative to the central values. A description of the uncertainties considered is provided below with their estimated values presented in the next section.

Pileup reweighting

Pileup reweighting is performed on the simulation to match the pileup profile between simulation and data [142, 143]. The pileup profile depends on the nominal pp inelastic cross-section for Run 2, which is measured to be 69.2 mb. The uncertainty on this cross-section is 4.6% which in turn is used to produce $\pm 1\sigma$ variations on the pileup profile. The shift in the JER for simulation is used to

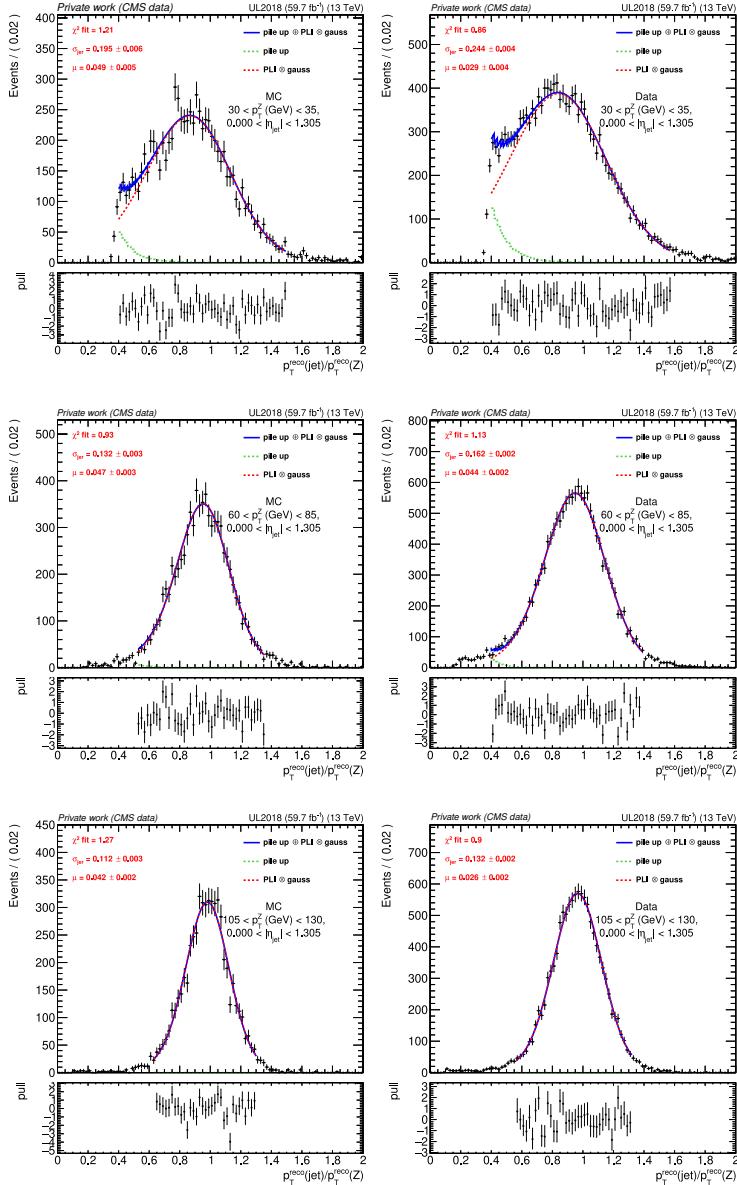


Figure 4.14: p_T balance response distribution for $30 < p_T^Z < 35$ GeV (top), $60 < p_T^Z < 85$ GeV (middle), $105 < p_T^Z < 130$ GeV (bottom) in the Barrel region for simulation (left) and data (right).

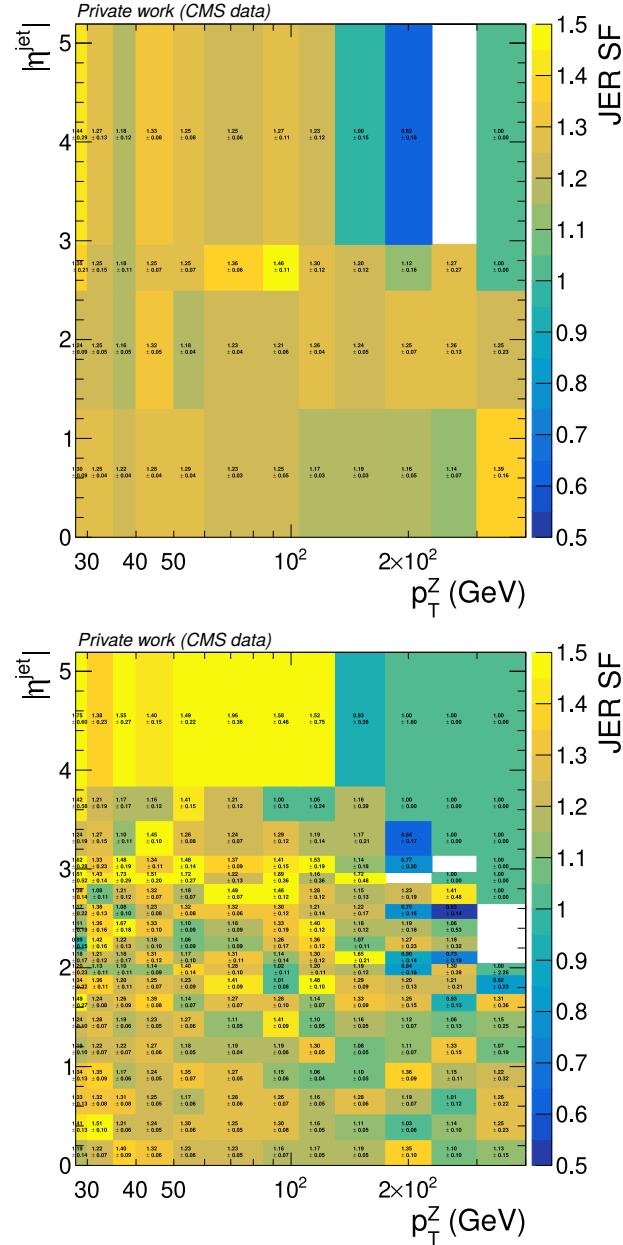


Figure 4.15: The Data/MC scale factor for the jet energy resolution as a function of $|\eta^{\text{jet}}|$ and p_T^Z using coarse (left) and fine (right) bins.

apply the uncertainty due to pileup reweighting.

Jet energy scale

The jet energy scale for data and simulation is corrected as described in Section 3.2.7. These calibrations come with associated uncertainties, which measure possible deviations of the simulation from the detector response observed in data. The impact of this source on the jet energy resolution is evaluated by applying the JES corrections with up and down variations to simulation. The resulting shift in the resolution is taken as the corresponding uncertainty.

Fit range

For each p_T balance distribution in simulation and data, the fit is performed within fixed bounds, excluding 2.5% of the distribution from the higher end (for all p_T^Z) and from the lower end (only for $p_T^Z > 50$ GeV). These bounds are chosen to suppress contributions from erratic tails. However, since these contributions depend on p_T and η , a systematic uncertainty is assigned to the choice of bounds. This uncertainty is evaluated by performing up and down variations, in which 3% and 2% of the distribution are excluded from each end, respectively.

Probe jet minimum p_T

A minimum p_T threshold of 12 GeV is applied on all jets considered. This requirement rejects any low p_T jets which may populate the p_T balance distributions, especially in the low p_T^Z bins. A systematic uncertainty for this selection is measured by varying the threshold to 10 GeV and 14 GeV and taking the resulting change in the resolution as the uncertainty.

Pileup template

The pileup template used for the fits is generated for simulation and data using the procedure described before. As an alternative, the non-gen. matched jets in simulation can be used to estimate the pileup contribution. The resolution is calculated by applying the p_T balance distribution for the non-gen. matched jets

as the pileup template and the resulting variation is taken as the uncertainty from this source.

Initial-state radiation and final-state radiation

Contributions from radiative processes such as initial-state radiation (ISR) and final-state radiation (FSR) play a significant role in the description of soft activity in simulation. In the low- p_T regime (below 100 GeV), the jet energy is particularly sensitive to the amount of soft activity in the event. The ISR and FSR variations are introduced by modifying the QCD coupling constant ($\alpha_S(\mu)$ where μ is the energy scale) in the parton shower. For ISR, this change is applied before the hard-scattering process, while for FSR it is applied after the hard-scattering process. Since these effects play a role at the particle-level, only the shift on the particle-level imbalance is applied to simulation and data while extracting the JER. The variations are performed by scaling μ by factors of 0.5 and 2.

4.5 p_T dependent JER results

The jet energy resolution scale factors shown in Fig. 4.15 provide the results in bins of p_T^Z and $|\eta^{\text{jet}}|$. To derive the scale factors as a function of jet kinematics only, a mapping between p_T^Z and p_T^{jet} is built using a linear fit to the profile of the p_T^{jet} distribution in bins of p_T^Z . The p_T dependent JER results as a function of p_T^{jet} are shown in Fig. 4.16 for the Barrel (top left), Endcap 1 (top right), Endcap 2 (bottom left) and Forward (bottom right) regions. Across all regions, the JER scale factor is typically found to be greater than 1.2, indicating that the jet energy resolution in data is larger than 20% for $p_T^{\text{jet}} < 100$ GeV. In the Barrel region, the scale factors increase as the p_T^{jet} decreases, reaching ~ 1.3 for $p_T^{\text{jet}} < 50$ GeV. The scale factors in the Endcap 1 and Forward regions show a much flatter trend, and can be approximated to a constant scale factor of 1.2. The complicated geometry of the Endcap 2 region leads to a more complex scale factor trend as a function of p_T^{jet} , reaching a maximum of > 1.4 .

The total uncertainty on the measurements, shown as green band in Fig. 4.16,

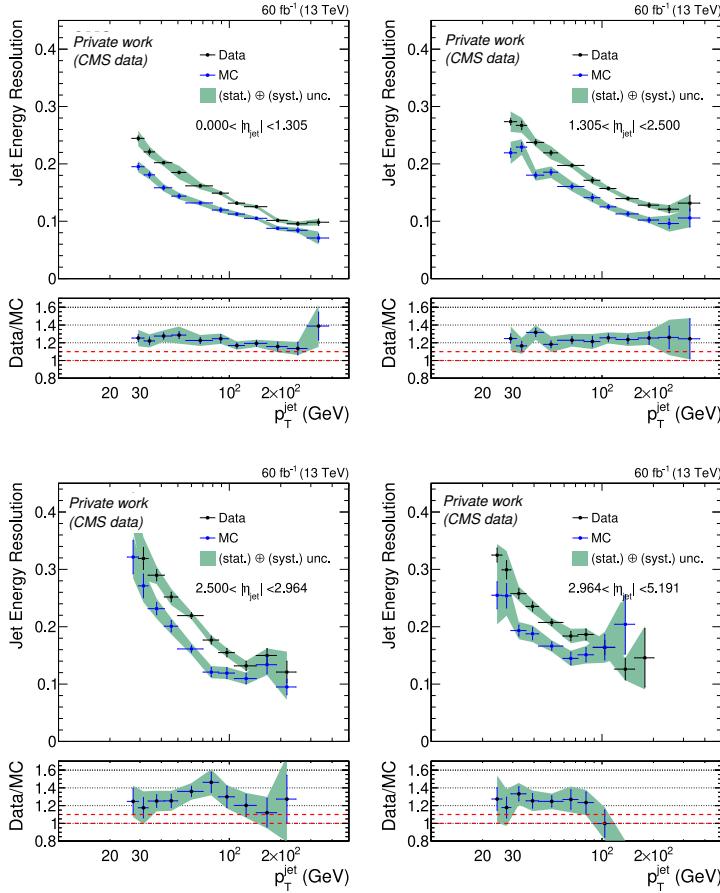


Figure 4.16: Jet Energy Resolution as a function of p_T^{jet} in the Barrel (top left), Endcap 1 (top right), Endcap 2 (bottom left) and Forward (bottom right) regions. The JER values are extracted from the fit using Eqn. 4.5.

consists of the statistical and systematic uncertainties described in the previous section. The relative uncertainty for the individual uncertainties on the JER for simulation (top left) and for data (top right) is shown in Fig. 4.17 for the Barrel region. These are propagated to the JER scale factors where the effect of low statistics is dominant at $p_T^{\text{jet}} > 100$ GeV. Below 100 GeV, a larger effect is observed with the pileup template variation in data. However, since the pileup contribution in simulation is different from data, this variation – introduced by

replacing the data template with non-gen. matched jets in simulation – serves as a conservative estimate. Among the other uncertainties, ISR modelling is observed to have one of the dominating effects on the JER measurement in simulation and data. This uncertainty is measured by its effect on the PLI, implying that a precise modelling of ISR in simulation would play an important role in improved measurements of the JER in both simulation and data.

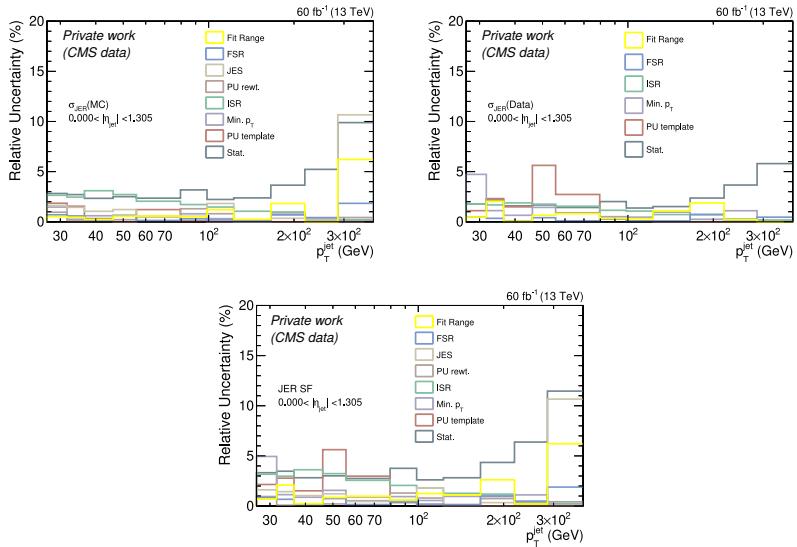


Figure 4.17: Relative uncertainty on the JER for simulation (top left), JER for data (top right) and the scale factor (bottom) in the Barrel region for different sources of uncertainty.

The previous Run 2 JER results (legacy datasets) suggest a scale factor of ~ 1.1 (Fig. 4.4). There was also no p_T dependence observed in the Barrel region with the previous results. In parallel to the results shown here using Z+jets events, efforts to probe the dijet event topology are currently ongoing in the CMS Collaboration to estimate the JER scale factors, with a focus on covering the $p_T^{\text{jet}} > 100$ GeV spectrum. The preliminary studies also suggest a p_T dependence of the JER scale factors and are found to be compatible in the overlapping region ($100 < p_T^{\text{jet}} < 150$ GeV) with the results in this thesis. Although the event selection is different between the two results, both rely on the same core philos-

ophy of using the p_T balance approach to derive the jet energy resolution. A new method based on the missing transverse momentum projection fraction (MPF) approach to jet energy corrections is also in development by the CMS Collaboration. The preliminary results with this method show a decreasing jet energy resolution scale factor below 100 GeV and therefore, contradict the results with the p_T balance method discussed so far. To understand the differences, this new methodology is explored in the context of the studies described so far and is presented in the next section.

4.6 Complementary methods: MPF method

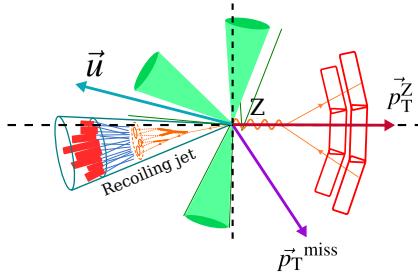


Figure 4.18: Schematic illustrating the components of the MPF method with $Z+jet$ events

The MPF approach relies on balancing all the hadronic activity in the transverse plane against the reference object p_T . Using $Z+jet$ events as a reference, the ideal topology for a hard-scattering event would consist of a Z boson recoiling against a single jet. However, the presence of radiative effects such as ISR and FSR causes additional hadronic activity which could leave some unclustered energy in the event. This can be interpreted as energy which is missing from the recoiling jet and would appear as the missing transverse momentum, p_T^{miss} . There will also be contribution from pileup to the net hadronic activity, denoted by the vector, \vec{u} (Eqn. 4.6). Balancing the net hadronic recoil against the Z boson leads to Eqn. 4.7. The vector \vec{u} can then be projected onto the Z -boson momentum

axis to study the response of the net hadronic activity, as shown in Eq. 4.8.

$$\vec{u} = \sum_{\text{jets}} \vec{p}_T \quad (4.6)$$

$$\vec{u} = -\vec{p}_T^{\text{miss}} - \vec{p}_T^{\text{ref.}} \quad (4.7)$$

$$R^{\text{MPF}} = \frac{\vec{u}_{\parallel}}{p_T^{\text{ref.}}} = 1 + \frac{\vec{p}_T^{\text{miss}} \cdot \vec{p}_T^{\text{ref.}}}{(p_T^{\text{ref.}})^2} \quad (4.8)$$

$$R^{\text{MPFX}} = \frac{\vec{u}_{\perp}}{p_T^{\text{ref.}}} \quad (4.9)$$

This is termed as the response of the MET projection fraction since R^{MPF} mentioned here is just the response of the MET projected onto the Z boson momentum axis and shifted by unity. Similar to the approach with p_T balance, the response distribution can then be used to measure the resolution of R^{MPF} . It would consist of the JER for the hard-scattering jet (σ_{JER}), the width of ISR and FSR response (σ_{ISR} , σ_{FSR}) and width of the pileup response (σ_{PU}). This can be parametrized as Eqn. 4.10 where the quantity of interest is σ_{JER} . To extract this, the MPF method is extended to measure the response in the orthogonal axis with respect to the Z boson momentum axis. The isotropic nature of pileup, ISR and to some extent, FSR, would contribute to the response of \vec{u} in this direction and can be written as Eqn. 4.11. Subtracting in quadrature the width of the response of \vec{u} in the parallel and orthogonal axis to the Z boson momentum, is then expected to cancel out all contributions except the jet energy resolution, leading to Eqn. 4.12.

$$\sigma_{\text{MPF}} = \sigma_{\text{JER}} \oplus \sigma_{\text{ISR}} \oplus \sigma_{\text{FSR}} \oplus \sigma_{\text{PU}} \quad (4.10)$$

$$\sigma_{\text{MPFX}} = \sigma_{\text{ISR}} \oplus \sigma_{\text{FSR}} \oplus \sigma_{\text{PU}} \quad (4.11)$$

$$\sigma_{\text{JER}} = \sigma_{\text{MPF}} \ominus \sigma_{\text{MPFX}} \quad (4.12)$$

It is important to note here that this would also remove any contribution coming from the pileup component present within the jet itself. This effect can be corrected for by adding back the pileup contribution using results from the Random Cone method [133]. However, this is a different study which is ongoing in the CMS Collaboration and falls beyond the scope of this thesis. The effect of this

is expected to be small and is neglected in the MPF studies presented here.

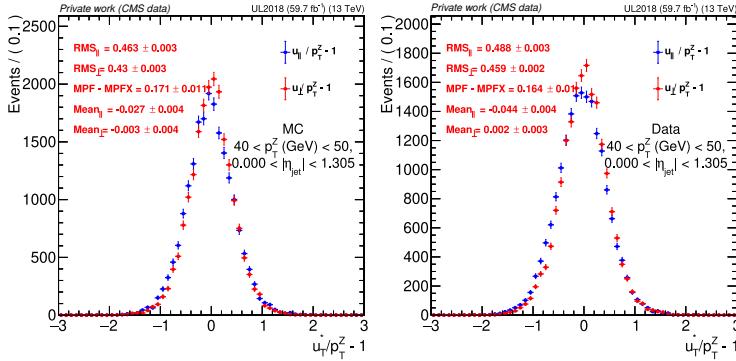


Figure 4.19: Response of parallel (blue) and perpendicular components (red) \vec{u} for $40 < p_T^Z < 50$ GeV in the Barrel region for simulation (left) and data (right).

The methodology discussed above is implemented to extract the JER after applying the same event selection which has been applied on the p_T balance results. The JER estimate is obtained by subtracting in quadrature the RMS of the response distribution for the parallel and the perpendicular components of \vec{u} in bins of p_T^Z and $|\eta_{jet}|$. An example bin is shown in Fig. 4.19 for $40 < p_T^Z < 50$ GeV in the Barrel region. As before, the mapping between p_T^Z and p_T^{jet} is used to study the p_T dependence of the JER in simulation and data and the corresponding scale factors in different regions of the detector.

Fig. 4.20 shows the p_T dependent results with the MPF method (left) in the Barrel region and is compared to the results with p_T balance method (right). For simulation, the jet energy resolution in both methods is seen to be compatible. This indicates that the same closure between the true JER and the extracted JER is expected for the MPF method applied here. The biggest differences come from the JER for data which are much smaller for the MPF method as compared to the p_T balance results. This results in JER scale factors of $\sim 1.05 - 1.1$. Additionally, the p_T dependence for the MPF method is different from what has been observed for the p_T balance. The MPF results suggest a decreasing trend for the JER scale factor as the p_T^{jet} decreases. The preliminary results from the groups in the CMS Collaboration working on the MPF method with dijet events indicate a similar trend. In comparison, for the p_T balance results presented in this

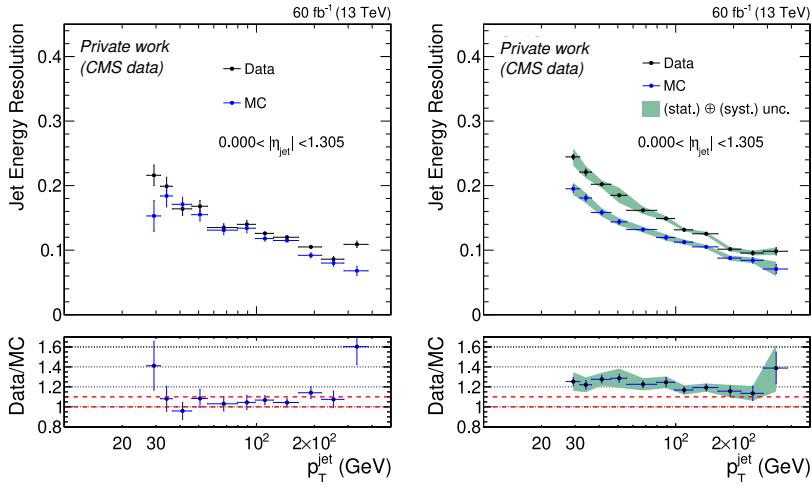


Figure 4.20: Jet Energy Resolution as a function of p_T^{jet} in the Barrel region with the MPF method (left) and the p_T balance method (right).

thesis, the opposite is observed where the scale factors increase with decreasing p_T^{jet} . While the reason for the fundamental differences in these results is under study in the CMS Collaboration, one of the major differences between both the results is the impact of the soft radiation and its modelling on the measurements. The MPF method is expected to be less biased by the description of both the pileup and ISR/FSR effects in simulation since these contributions are expected to cancel out in situ. For the p_T balance method, however, these effects have an influence on the extracted JER in data as it determines the PLI that is used in the fit. However, even with the inclusion of the ISR and FSR variations, the jet energy resolution for data has been observed to be still larger than what the MPF method predicts. As a result, further dedicated studies are needed to understand the root cause for this difference.

An alternative is to test both results by studying their influence on the physics explored in the CMS experiment. This would provide an indication of which result is closer to reality and serve as a validation. In the next section, a validation study using the JER scale factor results is presented and evaluating their potential impact on physics analyses.

4.7 Validation of JER scale factors

One of the key processes used to study QCD is the measurement of the inclusive jet cross section at the LHC. The ATLAS and CMS Collaborations have undertaken several such measurements at different center-of-mass energies. One of the central observables in these measurements is the jet p_T spectra which can act as a probe into parton momentum distributions in the proton, and consequently provide information on its structure. An accurate description of the jet p_T spectrum in simulation is therefore essential and requires all the calibration steps, including the corrections for the jet energy resolution.

Treating this as an inspiration, a validation study is carried out to apply the JER scale factors. For this study, the leading jet p_T spectrum in Z+jet events will be compared in simulation and data. The JER correction is applied to the simulation jets in a procedure called jet smearing which is described in the next section. The JER scale factors obtained from the p_T balance method and the MPF method presented before will be used separately. The post-smearing results which deviate the least from the distribution for data would indicate a more accurate description of the resolution in data.

4.7.1 Jet smearing

Two methods are commonly used to smear jets, termed as the scaling method and the stochastic smearing method. The recommended approach is a hybrid technique that combines both the methods.

The scaling method is applied to rescale the four-momentum of the jets which are matched to a particle-level jet. The matching criteria for a jet consists of a geometrical as well as p_T matching, as described in Eqn. 4.13. In the case of AK4 jets, a particle-level jet is required to be present in a $\Delta R < 0.2$ to pass the geometric matching criteria. Additionally, the particle-level jet p_T is required to be within $3\sigma_{\text{JER}}$ of the jet p_T where σ_{JER} is the true JER in simulation. Once a matching particle-level jet is found, the scaling factor, c_{JER} , is calculated

according to Eqn. 4.14

$$\Delta R < R_{\text{cone}}/2, |p_T - p_T^{\text{ptcl}}| < 3\sigma_{\text{JER}} p_T \quad (4.13)$$

$$c_{\text{JER}} = 1 + (s_{\text{JER}} - 1) \frac{p_T - p_T^{\text{ptcl}}}{p_T} \quad (4.14)$$

where s_{JER} is the JER scale factor. The scaling factor is also truncated at zero. Note that this method scales only the jets which are associated to particle-level jets with a well-reconstructed p_T and therefore, does not apply any scaling to non-gen. matched jets (i.e. pileup jets).

The stochastic smearing method provides an alternative approach by not depending on whether a jet is matched to a particle-level jet. The scaling factor is determined according to Eqn. 4.15

$$c_{\text{JER}} = 1 + \mathcal{N}(0, \sigma_{\text{JER}}) \sqrt{(\max(s_{\text{JER}}^2 - 1, 0))} \quad (4.15)$$

where the scale factor is s_{JER} and the true JER in simulation is σ_{JER} . The random term sampled from the normal distribution $\mathcal{N}(0, \sigma_{\text{JER}})$ is included to mimic the detector resolution spread. The recommended approach, the hybrid method, uses the scaling method for all jets which pass the matching criteria mentioned in Eqn. 4.13. For all other jets, mostly pileup, the stochastic smearing procedure is applied. In this study, only the scaling method has been applied as the leading jet p_T distribution is expected to be dominated by the hard-scattering jet.

For performing the jet smearing, the JER scale factors obtained for the p_T balance method and the MPF method are fitted with a linear function in each of the four coarse $|\eta^{\text{jet}}|$ regions as before. Figure 4.21 shows the JER scale factor distributions and their corresponding linear fits for both the methods in the Barrel region. As expected from the results seen before, two opposing trends are observed for the methods even after accounting for the $\pm 1\sigma$ variations on the fit. The JER scale factor values extracted from the fits are used for performing the jet smearing procedure.

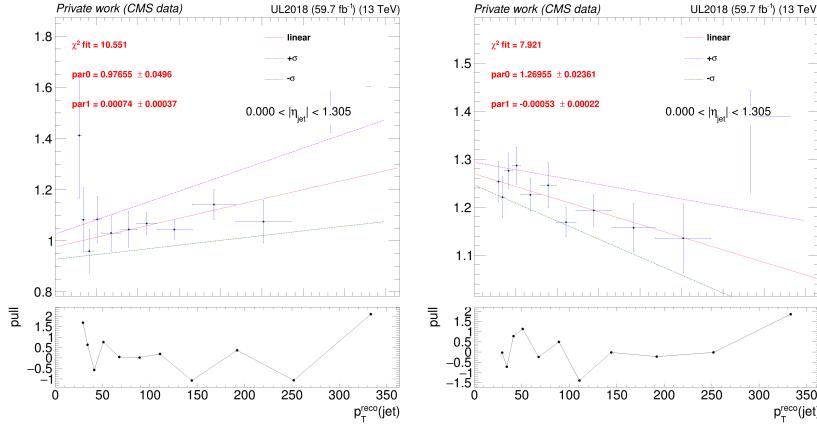


Figure 4.21: JER SF as a function of p_T^{jet} in the Barrel region with the MPF method (left) and the p_T balance method (right). A linear fit is performed on both the distributions in the range $30 < p_T^{\text{jet}} < 200$ GeV.

4.7.2 Datasets and Event selection

As mentioned in Section 4.2, the 2018 data set with $Z \rightarrow \mu\mu$ events are used to perform the validation studies. The events are selected with the dimuon trigger described before. For the Z +jet simulation samples, two NLO samples using MadGraph5_aMC@NLO and POWHEG are considered in addition to the LO MadGraph5 sample used in the previous studies. The details of these samples have been described in Section 4.2.

To study an inclusive p_T distribution of the leading jet, a very minimal set of cuts is included in the event selection. The jets follow the same reconstruction procedure and the calibration steps as described in Section 4.3, storing all jets with $p_T > 12$ GeV. The leptons are required to pass the same selection criteria to reconstruct a Z boson candidate with a dimuon mass of 70 to 110 GeV. To select the Z +jet event topology, the leading jet and Z boson candidate are required to be back-to-back by imposing the condition, $\Delta\phi(Z, \text{jet}) > 2.5$. No additional selection criteria have been placed on the jets in the event.

Since the objective of this study is to compare the impact of different JER corrections, the simulation should also be corrected for other major effects that contribute to differences between simulation and data. One source of data-

simulation discrepancy comes from the dilepton p_T distribution in simulation which can shift the data to simulation ratio by up to 20 – 30%. This effect is corrected by applying reweighting factors in bins of p_T^Z . This factor is calculated by normalizing the p_T^Z distribution in each simulation separately and reduces the data-simulation discrepancy to < 10%. Figure 4.22 shows the dilepton mass distribution for the simulation samples and data. There are additional corrections which can be applied to further improve the data simulation agreement but these have not been included here.

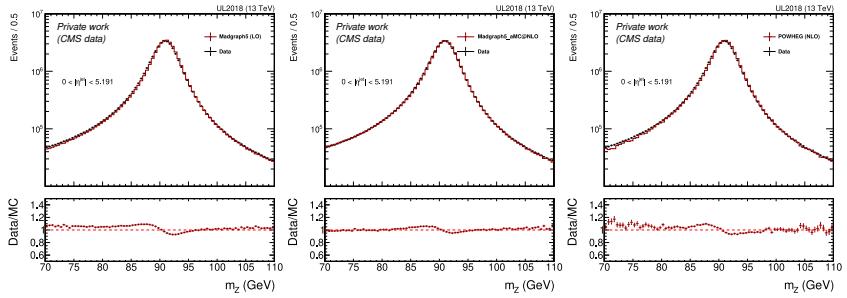


Figure 4.22: Reconstructed Z boson mass distribution for MadGraph5 LO (left), MadGraph5_aMC@NLO (middle) and Powheg NLO (right) samples in comparison with 2018 data.

4.7.3 Results

Figure 4.23 shows the leading jet p_T distribution for the different simulation samples in comparison to data. A similar jet p_T spectrum trend is present for each of the simulation samples. The Data/MC ratio for data and simulation with no JER correction applied (i.e. unsmeared MC) is $\sim 0.9 – 1$. With the JER correction from the p_T balance results (or smeared p_T balance) and the MPF results (or smeared MPF) included in the simulation, there is a small shift in the Data/MC ratio. A shift towards unity indicates better agreement of the data with the smeared simulation. In this case, the Data/MC results for smeared MPF samples show better Data-MC agreement as compared to the p_T balance results.

However, the difference between both the smeared distributions is observed to be small. To investigate further, the ratio of the MC samples smeared with p_T

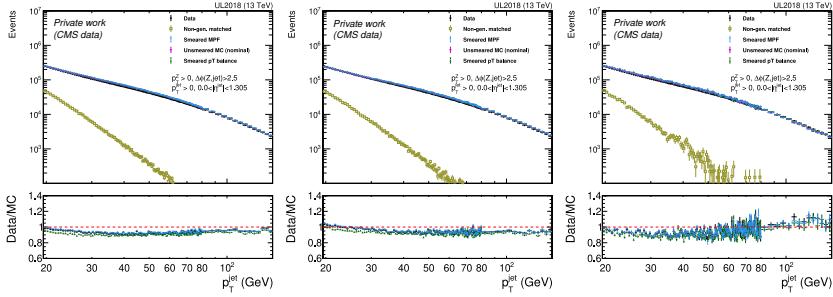


Figure 4.23: Leading jet p_T spectra in the Barrel region for MadGraph5 LO (left), MadGraph5_aMC@NLO (middle) and Powheg NLO (right) samples in comparison with 2018 data. The corresponding Data/MC ratios quantify the level of agreement between the simulation and data before and after the JER corrections are applied.

balance and MPF results with respect to the unsmeared samples are compared. In addition to this, the uncertainties coming from some common sources of uncertainty in the distributions of such observables are also considered. The jet energy scale uncertainty, ISR and FSR variations for the leading jet p_T spectrum of the unsmeared simulation have been included. The uncertainties on the JER scale factors are introduced by varying the fit (shown in Fig. 4.21) by $\pm 1\sigma$. The eventual variation in the spectrum is taken as the uncertainty due to the JER scale factors. In the results shown here, only the variation from p_T balance results is included as uncertainty for the JER scale factors.

Figure 4.24 shows the comparison between smeared samples and the unsmeared MC sample. It is observed that there is $< 5\%$ of difference between the smeared p_T balance results and the smeared MPF results. It is also important to note that the effect of smearing with either result is smaller than the uncertainty coming from the jet energy scale or ISR/FSR uncertainties. While the results shown here are for the Barrel region, similar observations are seen across the other detector regions, leading to two important inferences. The first is that the MPF method results seem to provide better Data-MC agreement as compared to the p_T balance results. The second is that the difference between the two smeared results are small and possibly covered by uncertainties coming from other sources such as the jet energy scale, ISR or FSR.

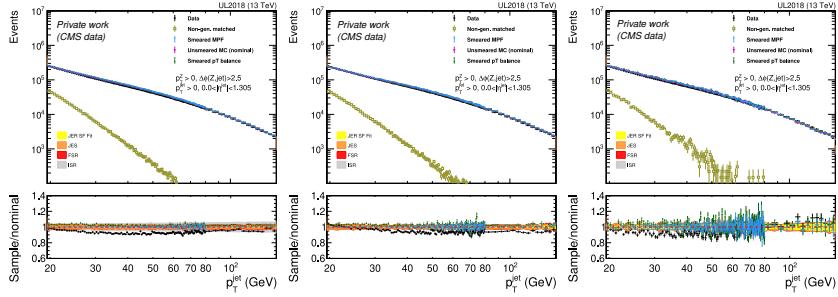


Figure 4.24: Leading jet p_T spectra in the Barrel region for MadGraph5 LO (left), MadGraph5_aMC@NLO (middle) and Powheg NLO (right) samples in comparison with 2018 data. The corresponding Sample/nominal ratios quantify the shift in the distribution after the JER corrections are applied. The systematic uncertainty is calculated for unsmeared nominal MC.

4.8 Chapter Summary

In this Chapter, a method of measuring the jet energy resolution was presented which extends the measurements to as low as $p_T^{\text{jet}} = 30$ GeV using a p_T balance approach and template-based fits. The resulting data-to-simulation scale factors are found to be of $\sim 1.1 - 1.3$. These results are compared with a newly developed method based on the MET projection fraction (MPF) approach, which yields scale factors of $\sim 1 - 1.1$. Finally, the impact of these two results is assessed with the leading jet p_T spectrum in $Z + \text{jet}$ events. The study suggests that the smearing from either method is small and that more extensive studies are required to fully understand the differences.

Chapter

5

Exclusive search for light scalar bosons in final-states with two muons and two hadrons

This chapter covers the details of the analysis searching for light scalar bosons in final states of two muons and two hadrons. This analysis has been the main focus of my thesis. As the main analyzer for this analysis, all results shown here have been performed by me unless stated otherwise.

5.1 Introduction

Multiple Standard Model extensions predict new particles with mass at or below the weak scale. In particular, light scalar bosons with masses of $\mathcal{O}(\text{GeV})$ appear in several SM extensions including cosmic inflation models [39–41], supersymmetry [42, 43], dark matter models [44, 45] and cosmological solutions to the hierarchy problem [46]. Chapter 1 discussed one such Standard Model extension that introduces a singlet scalar particle with $\mathcal{O}(\text{GeV})$ masses. The scalar boson mixes with the Higgs boson and inherits couplings to SM particles, allowing decays into visible final states. Depending on how small the mixing parameter is, the BSM scalar boson can decay close to the interaction point or

travel some distance in the detector before decaying. This leads to interesting and challenging scenarios that have been probed in various collider experiments and were briefly described in Chapter 1.2.3.

In this Chapter, a search for light BSM particles with masses below 2 GeV is presented with the CMS experiment. In this low-mass phase space, the hadronic decays of the BSM particles dominantly lead to only a hadron pair instead of a collimated jet of multiple hadrons. Based on the studies in Refs. [16–18], a SM extension with a light BSM scalar boson, S , in this mass range would lead to dominant hadronic decay modes through the exclusive channels, $S \rightarrow K^+K^-$ (for $m_S \geq 1.1$ GeV) and $S \rightarrow \pi^+\pi^-$ (for $m_S \leq 1.0$ GeV) (see Section 1.2.2, Fig. 1.4).

A search for low-mass scalar bosons in such exclusive decays was originally proposed in Ref. [144] for the LHCb experiment in fully-hadronic final states. The proposed strategy targets scalar bosons produced in exotic decays of the Higgs boson, $H \rightarrow SS$, followed by the scalar bosons decaying to charged kaon final states. In the CMS experiment, such a final state is hard to trigger on since it would be submerged by hadronic backgrounds. Therefore, an alternative strategy has been adopted in which one of the light scalar bosons is required to decay to muons which provides a signature that can be efficiently triggered on.

In the analysis presented in this thesis, a Higgs-mediated search for light scalar bosons, with final states of two oppositely charged muons and two oppositely charged kaons or pions ($pp \rightarrow H \rightarrow SS \rightarrow \mu^+\mu^-h^+h^-$) is performed. The complete final state, consisting of one pair of collimated hadrons (kaon or pion) and one of muons, is used to reconstruct the properties of the BSM light scalar. By comparing the masses, as calculated from the two track pairs, the displacement of the reconstructed decay vertices and by requiring that these particles are isolated, nearly all the background is suppressed. This allows for a unique sensitivity to prompt and displaced signatures. The analysis uses the Run 2 dataset collected by the CMS experiment. The mixing parameter between the Higgs boson and the BSM scalar is treated as a free parameter, targeting a search that is sensitive to a wide range of scalar boson lifetimes. Considering the values of the $S \rightarrow \mu^+\mu^-$, $S \rightarrow K^+K^-$ and $S \rightarrow \pi^+\pi^-$ branching fractions from Refs. [16, 17], the results of the search are interpreted as upper limits on the branching fraction of the Higgs boson decay to the scalar bosons as a function of the scalar mass

and lifetime.

The Chapter is structured as follows. Section 5.2 discusses the signal modelling and simulation followed by a description of the datasets used to study the backgrounds and perform the search. In Section 5.3, the trigger strategy for the analysis is presented. The complete event selection is then described in Section 5.4, starting from the object selection to the final handles used to reduce the backgrounds. Section 5.5 presents the background estimation method followed by the results of the event selection analysis in Section 5.6. The sources of uncertainty and the relevant corrections applied to simulation are discussed in Section 5.7. Finally, the statistical interpretation and the sensitivity of the search are presented in Section 5.8.

5.2 Signal simulation and datasets

5.2.1 Signal simulation

The signal benchmark model considered for the analysis is the Higgs boson mediated production of BSM scalar bosons, S . The Higgs boson can be produced at the LHC through various processes. Some of the major production modes and their corresponding cross-sections are listed in Table 5.1. The leading production mode is the gluon-gluon fusion process which accounts for nearly $\sim 87\%$ of the total Higgs boson production cross-section [60]. Based on this information, the signal simulation for the benchmark model uses the gluon-gluon fusion production mode (ggH) for the Higgs boson, followed by its decay to two BSM scalar bosons which in turn decays as shown in Fig. 5.1. In the later stages of the analysis, the Higgs boson production modes with vector boson fusion and with associated vector boson production, have also been included.

The first part of the signal simulation is the production of SM Higgs. The Matrix Element calculator, `POWHEG2.0` [145–148], is used to generate the ggH process at next-to-leading order (NLO) accuracy. This provides the parton level information for the hard-scattering event. The next part of the simulation consists of defining the rest of the hard-scattering event which contains the exotic decay of the SM Higgs boson to the BSM scalar particles and the targeted final state

Production mode	Cross-section (pb)	Process
Total	55.6	$pp \rightarrow H$
Gluon-gluon fusion (ggH)	48.6	$gg \rightarrow H$
Vector boson fusion (VBF)	3.78	$qq \rightarrow q\bar{q}H$
Associated production with vector boson (VH)	2.26	$q\bar{q} \rightarrow VH$ $V=W^\pm, Z$

Table 5.1: Some major production modes for the SM Higgs boson and their cross-sections at $\sqrt{s} = 13$ TeV, evaluated at NNLO accuracy in the QCD coupling. Numbers taken from Ref. [60].

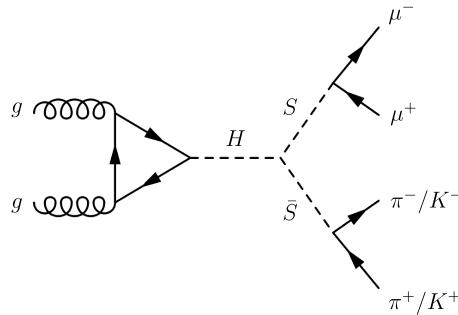


Figure 5.1: Feynman diagram of Higgs boson mediated BSM light scalar boson production in gluon-gluon fusion processes. The diagram illustrates the targeted decay signature of a pair of muons and a pair of hadrons.

particles. Along with the parton shower and hadronization process, this step is modelled with PYTHIA8.240 [103] using the CP2 underlying event tune [149].

To save computational power, only the targeted event topology for this analysis is generated in the signal samples. Accordingly, even though the SM Higgs boson can decay through several channels, it is forced to decay only to two BSM scalar particles with a defined mass, m_S , and proper lifetime, $c\tau$.

The targeted BSM scalar bosons in this search have a mass between 0.4 to 2 GeV with a defined mean proper decay length, $c\tau$, which determines how far the particle travels in the detector. The mean decay length (or travel distance) can be calculated as $L = \beta\gamma c\tau$, where $\beta\gamma$ is the Lorentz boost of the particle. In this search, the targeted final states, muons and hadrons, are identified using

the information from the tracker. Accordingly, any BSM scalar boson which decays beyond the tracker volume falls outside the scope of this analysis. In Section 3.2.1, the capabilities of the CMS tracker were discussed where Fig. 3.2 shows that the tracking efficiency falls below 10% as one moves radially away by 600 mm from the interaction point. This restricts the decay length up to which the BSM scalar boson can be probed. Assuming a Higgs boson produced at rest, momentum conservation and the relation, $\beta\gamma = p/m$, provide a rough estimate of the mean decay length. This in turn guides the range of the $c\tau$ that this analysis is sensitive to. As a reference, for $m_S = 2$ GeV and $c\tau = 100$ mm, the mean decay length turns out to be 3125 mm. For smaller masses, this is larger and indicates that for BSM scalar bosons with $c\tau = 100$ mm, a substantial number of the decays already start to occur beyond the tracker volume. This provides a range in $c\tau$ to which the analysis is expected to be sensitive. Accordingly, four points for the proper decay length are chosen ($c\tau \in [0.1, 1, 10, 100]$ mm) for each BSM scalar mass, m_S . Additionally, the events which lie beyond the tracker volume are removed from generation. The impact of this cut is discussed later in this Section 5.2.1.

Once the properties of the BSM scalar boson are specified, the next step is simulating the decay modes for the scalar. One of the BSM scalar bosons is fixed to decay to a pair of muons. The other BSM scalar boson is set to decay to a pair of charged hadrons. The BSM scalar bosons with masses below 1.1 GeV cannot decay into kaons due to kinematic constraints. Accordingly, the BSM scalar boson is set to decay to a pair of charged pions only. For $m_S \geq 1.1$ GeV, the dominant hadronic decay mode is into charged kaons. There are also contributions from decays to charged pions and four-pion final states. However, since the contributions from these decay channels are small, only the charged kaons final state is considered. This was confirmed by generating a separate signal simulation for $m_S = 2$ GeV with both the charged pion and the charged kaon decay modes included. This is discussed in Section 5.4.2 during the generator-level studies of the signal kinematics.

The complete set of signal masses considered for the signal samples is shown in Table 5.3. For the purpose of illustration, a reference branching fraction of $\mathcal{B}(H \rightarrow S\bar{S}) = 1\%$ is assumed. Using the ggH cross-section of Table 5.1, the

m_S (GeV)	$\mathcal{B}(S \rightarrow \mu^+ \mu^-)$	$\mathcal{B}(S \rightarrow \pi^+ \pi^-)$	m_S (GeV)	$\mathcal{B}(S \rightarrow \mu^+ \mu^-)$	$\mathcal{B}(S \rightarrow K^+ K^-)$
0.4	0.144	0.571	1.1	0.04	0.307
0.6	0.095	0.603	1.2	0.071	0.401
0.8	0.061	0.626	1.4	0.085	0.413
0.9	0.034	0.644	1.6	0.117	0.412
1.0	0.01	0.654	1.8	0.181	0.377
			2.0	0.239	0.35

Table 5.2: Branching fractions for the BSM scalar decays to a muon pair and to a pion or kaon pair [17], in the targeted range of BSM scalar masses (m_S). For $m_S \leq 1$ GeV, the hadronic decay mode to kaons is not allowed because of kinematic constraints. The charged kaon decay mode dominates over the charged pion decay mode for $m_S \geq 1.1$ GeV.

branching ratios mentioned in Table 5.2, and the full Run 2 luminosity of 138 fb^{-1} , the expected number of signal events for each mass hypothesis is shown in Table 5.3.

m_S (GeV)	Expected events in Run 2	m_S (GeV)	Expected events in Run 2
0.4	5523	1.1	824
0.6	3848	1.2	1912
0.8	3335	1.4	2358
0.9	1470	1.6	3238
1.0	439	1.8	4584
		2.0	5619

Table 5.3: Mass hypotheses considered for the signal simulation and the number of expected events in Run 2 ($\mathcal{L} = 138 \text{ fb}^{-1}$), based on the theoretical cross-section of ggH process, reference of $\mathcal{B}(H \rightarrow SS) = 1\%$, and the corresponding branching fractions $\mathcal{B}(S \rightarrow \mu^+ \mu^-)$ and $\mathcal{B}(S \rightarrow K^+ K^-)$ or $\mathcal{B}(S \rightarrow \pi^+ \pi^-)$ based on m_S .

Lifetime Reweighting

For the signal sample generation, a range of mass hypotheses is considered to ensure a good granularity in the range $0.4 < m_S < 2$ GeV. For the granularity in lifetime, four samples for each m_S have been simulated: $c\tau \in [0.1, 1, 10, 100]$ mm. To ensure a good scan of BSM scalar bosons with $c\tau \leq 100$ mm, the exponential nature of a particle's decay is used to generate additional samples with intermediate lifetimes.

Considering natural units (i.e. $c = 1$), the particle decay length follows the exponential distribution with the decay constant defined as τ . Since the 2 BSM scalar bosons are independently decaying, the distribution of signal events can be described by the product of two long-lived particle (LLP) decay probabilities as shown in Eqn. 5.1

$$p(t_1, t_2 | \tau) = \frac{1}{\tau^2} \exp[-t_1/\tau] \exp[-t_2/\tau] \quad (5.1)$$

$$t_i = \frac{\text{LLP travel distance (lab frame)}}{\beta_i \gamma_i} \quad (5.2)$$

where t_1, t_2 denotes the life-time of the first LLP and the second LLP in their rest frame respectively. These can be defined as Eqn. 5.2 by measuring the LLP travel distance at the generator-level in the lab frame and using the Lorentz boost of the particle ($\beta \gamma$).

Equation 5.1 can be used to define the probability distribution with a known $\tau = \tau_{old}$ and for a new $\tau = \tau_{new}$. On taking the ratio of the equation for τ_{old} and τ_{new} , a weight is obtained which can be assigned as the ‘LLP reweighting’ factor to the sample with τ_{old} to generate the sample with τ_{new} life-time.

$$w = \left(\frac{\tau_{old}}{\tau_{new}} \right)^2 \exp[(t_1 + t_2) \left(\frac{1}{\tau_{old}} - \frac{1}{\tau_{new}} \right)] \quad (5.3)$$

Using the formulation in Eqn. 5.3, the intermediate lifetimes can be generated by reweighting the existing samples to either larger lifetimes or lower lifetimes. Additionally, a particular lifetime sample could also be produced by using a mix of the next larger simulated sample and the next smaller simulated sample as well. In this analysis, the additional lifetimes are generated by reweighting to lower lifetimes from the existing samples because of statistical limitations that arise when reweighting to higher lifetimes. Fig. 5.2 demonstrates the lifetime reweighting procedure as a proof of concept by comparing the travel distance (L) of a BSM scalar boson of the simulated samples and their reweighted counterparts. It is seen that the reweighting gives the same distribution after reweighting of a larger lifetime to a known smaller lifetime sample, therefore validating the procedure. In the end, a good granularity in $0.01 < c\tau < 100$ mm is assured with the reweighting procedure described here.

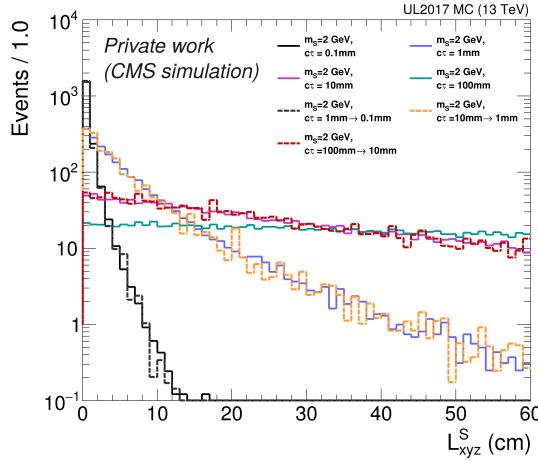


Figure 5.2: Demonstration of lifetime reweighting using the distance travelled by the BSM scalar boson at generator level for $m_S = 2$ GeV. Reweighting is performed from existing samples to smaller lifetimes. The resulting reweighted sample is then compared with nominal samples generated at the corresponding $c\tau$.

Generator-level filter

As this search targets BSM scalar bosons that decay within the tracker volume, a filter is applied during the sample generation to only keep events which will be reconstructed and can be detected in the CMS experiment. This filter removes events in which either of the BSM scalar bosons decays at $L_{xy} > 600$ mm, where L_{xy} denotes the transverse decay length of the scalar in the lab frame. Additionally, a $p_T^\mu > 5$ GeV cut is applied on each of the final state muons to reject events with soft muons since these are typically not well-reconstructed. The impact of the filter is described below.

The efficiency of the filter is measured as a function of the lifetime for $m_S = 0.8$ GeV and $m_S = 1.2$ GeV by generating samples at small intervals in $c\tau$ from $c\tau = 0.01$ to 100 mm. Fig. 5.3 shows the filter efficiency for the two considered mass points. At lower lifetimes, the filter efficiency is flat at $\approx 80\%$ indicating that the filter rejects events based on the p_T threshold applied on the muons. For the larger mass point ($m_S = 1.2$ GeV), it is seen that the drop appears at a larger

lifetime (~ 8 mm) as compared $m_S = 0.8$ GeV (~ 3 mm). This is expected since a smaller mass point has a larger decay length in the lab frame for the same proper decay length as a result of the larger boost.

A sigmoid-type fit is performed on it to extract an approximate estimate of the efficiency for all lifetimes between 0.01–100 mm. The exact functional form used is specified in Eqn. 5.4.

$$f(x) = \frac{p_0}{1 + p_1 x^{-p_2}} + p_3 \quad (5.4)$$

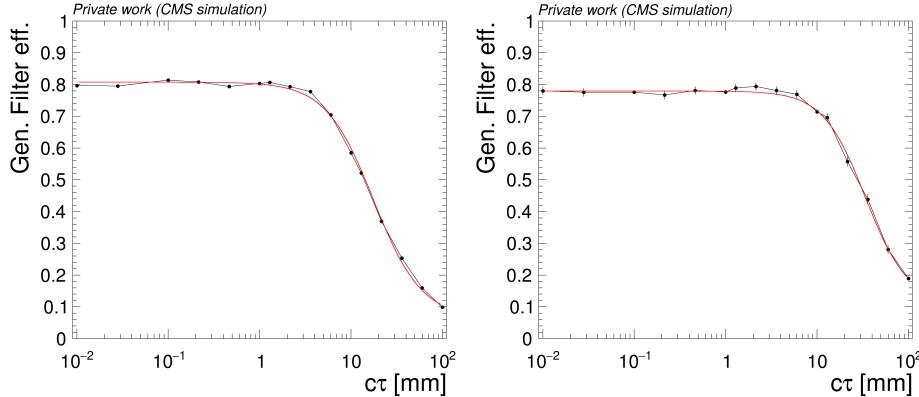


Figure 5.3: Generator-level filter efficiency for $m_S = 0.8$ GeV (left) and $m_S = 1.2$ GeV (right) as function of scalar lifetime. The red line denotes a sigmoid-type fit using Eqn. 5.4. As the lifetime increases, larger number of events are rejected by the filter since at least one of the scalars has $L_{xy} > 600$ mm.

m_S	\mathcal{E}_{filter}			
	$c\tau = 0.1$	$c\tau = 1$	$c\tau = 10$	$c\tau = 100$ mm
0.8 GeV	0.807	0.801	0.599	0.105
1.2 GeV	0.779	0.778	0.717	0.189

Table 5.4: Generator filter efficiency for $c\tau = 0.1, 1, 10, 100$ mm with samples of $m_S = 0.8, 1.2$ GeV.

The efficiency of the filter is an important factor which needs to be considered when estimating the final limits of the analysis, shown in Section 5.8. For

the samples with charged kaon final states, the efficiency values from the fit obtained for $m_S = 1.2$ GeV will be used (i.e. $m_S \geq 1.1$ GeV). The fit obtained for $m_S = 0.8$ GeV will be used for the masses $m_S \leq 1$ GeV where the charged pion final states are present. Table 5.4 shows the efficiency for the produced lifetime samples.

Higgs p_T correction

In the signal samples considered for the analysis, the generation of the ggH process is carried out at the next-to-leading order (NLO). Some recent theoretical calculations [150] provide the simulation of this process at the next-to-next-to-leading order accuracy, including the matching to a parton shower (NNLOPS). One of the significant differences that arise between the two simulations is the description of the Higgs boson p_T spectrum. This is illustrated in Fig. 5.4 where the generated signal samples in this analysis (blue) and a similar CMS analysis [61] using samples at NLO accuracy are compared with the NNLOPS results from Ref. [150]. It is observed that the distribution for the signal simulation differs by up to $\sim 20 - 30\%$ in the bulk of the distribution. To account for the impact of the NNLOPS results on the signal samples, the ratio (red) shown in the bottom plot of Fig. 5.4 is applied as a reweighting factor to the events in the signal samples. This leads to an increase of the signal yield by a few percent.

5.2.2 Data and background simulation

The data used in this analysis are from proton-proton collisions recorded at $\sqrt{s}=13$ TeV in 2016 (early, late), 2017 and 2018 corresponding to 36.3 (19.5, 16.8) fb^{-1} , 41.5 fb^{-1} , and 59.8 fb^{-1} , respectively. The 2016 dataset is split into two periods because the early 2016 data-taking was affected by saturation effects in the silicon strip tracker, which were fixed for the late 2016 data taking (see Section 2.2.2). Each era is individually processed and analyzed to account for the differences in data-taking conditions, such as pileup. These effects are reflected by simultaneously using simulation samples, corrections and systematic uncertainties for each era separately. The samples for this analysis have been generated in the final and the most refined reprocessing campaign for Run 2, referred to as

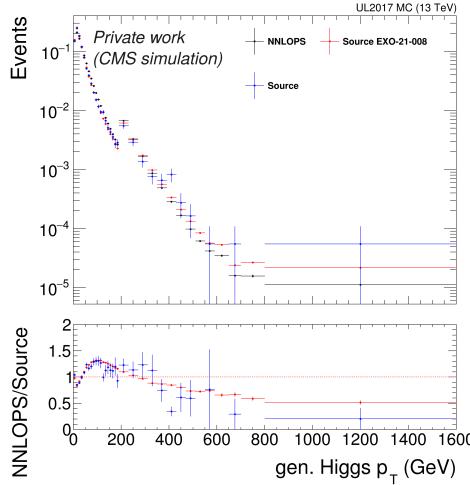


Figure 5.4: The Higgs boson p_T distribution for the produced signal sample (referred to as Source) (blue) (red) (blue), a reference CMS analysis [61] (red), and the results for the NNLOPS simulations from Ref. [150] is shown. The ratio between NNLOPS simulations and the other simulations considered here is shown in the bottom subplot.

the ultra-legacy (UL) samples. Accordingly, the dataset names in the figures in this Chapter are attached with the ‘UL’ prefix.

To study the major background processes, all the simulation samples mentioned in Table 5.5 have been used, of which quantum chromodynamics (QCD) multijets is the major contributor followed by $t\bar{t}$ processes. A muon-enriched QCD sample (at least one muon with $p_T > 5$ GeV) has been used in the analysis to understand the QCD background containing muon final states. The samples are generated with various matrix element generators, followed by parton shower and hadronisation processes modelled with PYTHIA8.240 [103], using the CP5 underlying event tune [140]. The simulation samples are weighted according to the cross-sections mentioned and scaled to the luminosity for comparison to data. The complete set of background samples were generated for the 2017 data set only for developing the analysis strategy and the event selection. For the early 2016, late 2016 and 2018 eras, only the dominant backgrounds, muon-enriched QCD and $t\bar{t}$ samples, were used since the contribution from other

samples was negligible.

In the following sections (Section 5.3– 5.5), the figures use 2017 data and corresponding simulation samples to support the event selection strategy and the background estimation method, unless stated otherwise. The same is applied for all the other eras, leading to a result on the full Run 2 dataset, shown in the figures of Section 5.6– 5.8.

Process	Additional information	ME Generator	σ (pb)
QCD (muon-enriched)	$15 < p_T^{\text{hard}} < 20 \text{ GeV}$	PYTHIA8	2797000
	$20 < p_T^{\text{hard}} < 30 \text{ GeV}$		2518000
	$30 < p_T^{\text{hard}} < 50 \text{ GeV}$		1361000
	$50 < p_T^{\text{hard}} < 80 \text{ GeV}$		377800
	$80 < p_T^{\text{hard}} < 120 \text{ GeV}$		87740
	$120 < p_T^{\text{hard}} < 170 \text{ GeV}$		21070
	$170 < p_T^{\text{hard}} < 300 \text{ GeV}$		7019
	$300 < p_T^{\text{hard}} < 470 \text{ GeV}$		622
	$470 < p_T^{\text{hard}} < 600 \text{ GeV}$		58
	$600 < p_T^{\text{hard}} < 800 \text{ GeV}$		18
$t\bar{t}$	$800 < p_T^{\text{hard}} < 1000 \text{ GeV}$	POWHEG	3.3
	$p_T^{\text{hard}} > 1000 \text{ GeV}$		1.6
$t\bar{t}$	Fully hadronic	POWHEG	377.9
	Semi-leptonic		366.9
	Fully leptonic		89.1
Single top	tW / tZq / tqb / tb	POWHEG / Madgraph5_aMC@NLO	251.6
W+jets	$W^\pm \rightarrow l^\pm \nu$	MadGraph5	53940
Z+jets ($Z \rightarrow l^+l^-$)	$0 < p_T(Z) < 50$	MadGraph5_aMC@NLO	5352
	$50 < p_T(Z) < 100 \text{ GeV}$		398
	$100 < p_T(Z) < 250 \text{ GeV}$		93
	$250 < p_T(Z) < 400 \text{ GeV}$		3.6
	$400 < p_T(Z) < 650 \text{ GeV}$		0.5
	$p_T(Z) > 650 \text{ GeV}$		0.05
$t\bar{t}V$	$t\bar{t}Z$	Madgraph5_aMC@NLO	0.7
	$t\bar{t}\gamma$	Madgraph5	1.5
	$t\bar{t}W$	POWHEG	0.6
	$t\bar{t}H$	POWHEG	1.1
VV	WW / WZ / ZZ / W γ / Z γ	POWHEG / Madgraph5_aMC@NLO	551.9
VVV	WWW / WWZ / WZZ / ZZZ	Madgraph5_aMC@NLO	0.46

Table 5.5: Samples used as background simulation with 2017 data-taking conditions and cross-sections; For other years, only QCD & $t\bar{t}$ samples with respective data-taking conditions have been generated.

5.3 Trigger Selection

The first stage of forming the analysis strategy, is finding the optimal trigger. The final state particles muons are coming from the decays of the BSM scalars ($m_S < 2$ GeV) which, in turn, come from the SM Higgs boson with mass $m_H = 125$ GeV. Due to the large mass difference between the Higgs boson and the BSM scalar, the BSM scalars are generally produced with a high p_T . This in turn ensures that at least one of the decay products from the BSM scalar is produced with a relatively high p_T which can be used to trigger the signal event. This is shown in Fig. 5.5 (left) where the mean p_T of the leading muon (the muon with the highest p_T) at the generator level (or the leading muon gen. p_T) is on average at around 40 GeV while a majority of the subleading muons have a generator level (hereafter abbreviated to gen.) p_T of less than 15 GeV.¹ Furthermore, the angular separation between the two muons, shown by the ΔR distribution in Fig. 5.5 (right), indicates that the muons are very close-by (or collimated). As a comparison, for dimuon decays of inclusively produced Z boson, the most probable value of ΔR between the muons is around 3 (back-to-back decays).

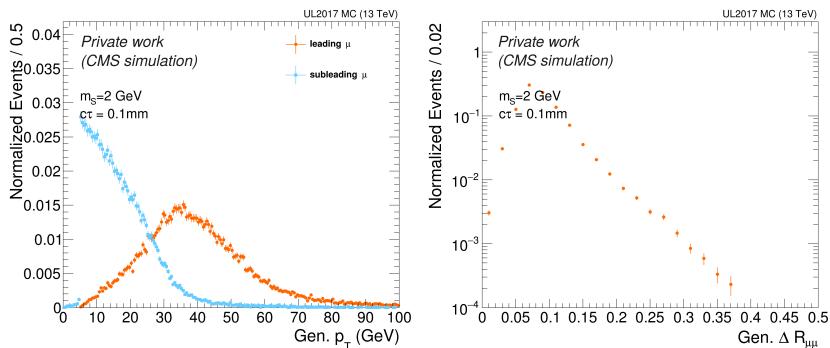


Figure 5.5: Gen. p_T for leading and subleading muon (left) and ΔR between leading and subleading muon (right) for $m_S = 2$ GeV, $c\tau = 0.1$ mm.

Given these properties of the muons in the signal events, various trigger op-

¹The difference between the p_T of the two muons follows from Lorentz transformation of the rest frame of the scalar (where the daughter particles have same momentum) to the lab frame (where the p_T of the two particles become asymmetric primarily due to the presence of the lorentz factor γ in the momentum calculations).

tions in Run 2 were tested. For the double muon triggers, several options required a relatively high p_T threshold on the subleading muon. In cases where the p_T threshold was not an issue, a requirement on the amount of track activity around the objects, quantified as isolation, was placed to keep the trigger rate low. However, due to the collimation between the leading and subleading muon, the isolation for the subleading muon includes the p_T contribution of the leading muon which causes it to fail the requirement.

Eventually, the most efficient trigger for the analysis was found to be the isolated single muon trigger across all years of Run 2 (with p_T threshold at > 24 GeV for 2016, 2018 and at > 27 GeV for 2017). The signal efficiency for the trigger as a function of the leading muon gen. p_T is shown in Fig. 5.6 (left) for a signal mass hypothesis of $m_S = 2$ GeV and multiple $c\tau$. The trigger is observed to be highly efficient for the signal samples with small lifetimes ($c\tau \leq 1$ mm) but drops significantly at larger lifetimes. This is because the isolated single muon trigger applies a maximum cut on the transverse impact parameter (d_{xy}) of the trigger muon. This is seen in Fig. 5.7 (left) where the signal efficiency drops to zero beyond a fixed d_{xy} value of the leading muon.

In 2018, two displaced dimuon triggers help in recovering some of the signal efficiency for signal hypotheses with $c\tau > 1$ mm. These triggers have been extensively used in Ref. [122] and are designed for capturing events with displaced muons. These triggers require at least two muons reconstructed in the muons systems only with $p_T > 23$ GeV each. One of these triggers is more tuned towards capturing events with muons that are seeded with track segments pointing towards the primary vertex while the other removes this constraint. As the triggers do not impose any requirements on the isolation or the d_{xy} of the trigger muon, they help in increasing the sensitivity to larger lifetimes, as shown by the combined trigger efficiency in Fig. 5.7 (right). However, they are limited by the p_T threshold on the subleading muon and hence, the isolated single muon trigger still contributes the most to the signal efficiency, as shown in Fig. 5.6 (right). The complete list of triggers used in this analysis is listed in Table 5.6.

The trigger efficiency shown in Fig. 5.6 and 5.7 however only illustrates the expected efficiency in signal simulation, and does not correspond to the efficiency expected in data. To account for the systematics differences between the

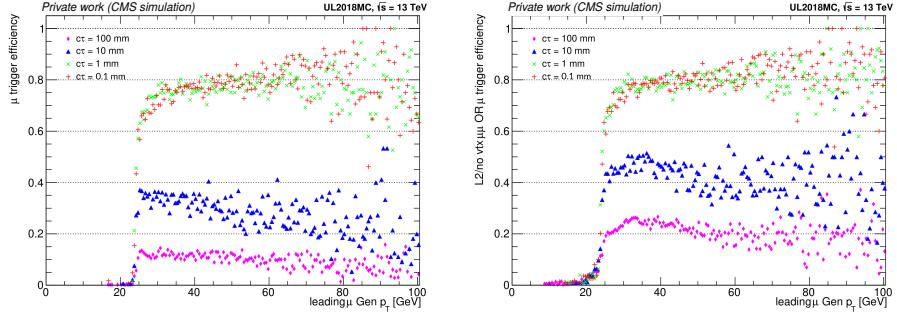


Figure 5.6: Signal efficiency for only isolated single muon trigger (left) and combination of displaced dimuon triggers with isolated single muon trigger (right) for UL2018 MC samples with scalar of mass 2 GeV as a function of leading gen. muon p_T . Some signal events are recovered for $c\tau = 10, 100$ mm samples with the displaced dimuon triggers.

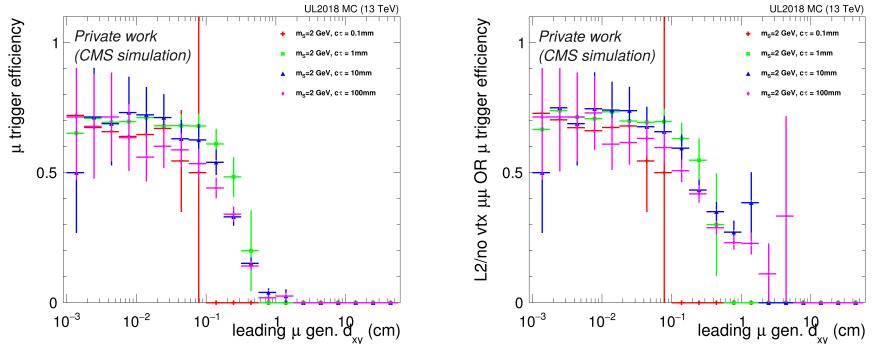


Figure 5.7: Signal efficiency for only isolated single muon trigger (left) and combination of displaced dimuon triggers with isolated single muon trigger (right) for UL2018 MC samples with a BSM scalar boson of mass 2 GeV as a function of leading gen. muon d_{xy} (transverse impact parameter). The displaced dimuon triggers lead to acceptance of signal events at larger displacement.

performance of the trigger in data and simulation, a separate study on measuring the efficiency for the isolated single muon trigger is carried out. The corresponding data-MC scale factors of this study provide the corrections that are then propagated to simulation. A dedicated study to derive these corrections was performed and is discussed later in Section 5.7.

Year	Trigger path
2016	HLT_IsoMu24
2017	HLT_IsoMu27
2018	HLT_IsoMu24 OR HLT_DoubleL2Mu23NoVtx_2Cha OR HLT_DoubleL2Mu23NoVtx_2Cha_CosmicSeed

Table 5.6: The list of triggers used in the analysis.

5.4 Event selection

This section presents the event selection criteria which are designed to enhance the sensitivity to signal-like events and simultaneously suppress the background contribution. The section is divided into several subsections, each focusing on different aspects of the selection strategy. Section 5.4.1 describes the quality cuts imposed on the physics objects used in the analysis. Section 5.4.2 outlines the analysis strategy, based on the signal sample kinematics. This is followed by the sections which detail the various selection handles.

5.4.1 Object Selection

Muons

The muon reconstruction described in Section 3.2.5 provides the base set of muons on which additional requirements are included to form the muon collection that is used in the analysis. The muons are required to pass the loose ID and are required to have $p_T > 5$ GeV and $|\eta| < 2.4$. The $|\eta| < 2.4$ and the loose ID requirement ensure that the muons lie within the tracker acceptance region and have a track associated to them. This is vital for reconstructing the secondary vertex where the muons in signal-like events originate from. Applying a minimum p_T requirement helps suppress the contribution from background muons.

Additional p_T requirements are imposed according to the trigger for the event. Since the muon p_T reconstructed at the online trigger level is different from the muon p_T after the full offline reconstruction, the trigger efficiency for the events with a muon near the trigger threshold (with rapidly rising efficiency

as a function of the offline muon p_T , referred to as the turn-on curve) are difficult to model and this may lead to distorted kinematic distributions. To address this, additional p_T cuts are chosen such that the events have muons which lie in the stable plateau region of the trigger efficiency. For the events selected with the isolated single muon trigger, the leading muon in the collection, which is most likely the muon that triggers the event, is required to pass the following offline selection criteria: for the 2016 and 2018 data set, $p_T > 26$ GeV is imposed and for the 2017 data set, $p_T > 29$ GeV is required on the leading muon. For the events in the 2018 data set that are triggered by the displaced dimuon triggers, at least two muons with $p_T > 26$ GeV are required.

Charged Hadrons

The charged hadron collection consists of charged particle tracks from the PF algorithm (presented in Section 3.2.3), which identifies the track as a charged pion by default. Depending on the targeted final state hadron, charged pion or charged kaon, the particle four-momenta is redefined. After the redefinition is carried out, the particle is required to have $p_T > 5$ GeV and $|\eta| < 2.4$. A similar reasoning as for the muon is applicable here to ensure that the large rate of background charged hadron tracks is suppressed.

Secondary Vertex

The final state consists of two charged hadrons and two muons, each originating from the decay of the respective BSM scalar boson. By identifying the correct pair of charged hadrons or muons, the properties of the originating BSM scalar boson can be reconstructed. One of these properties is the decay length which is measured as the distance of the decay vertex (or secondary vertex) from the primary vertex in the event. Since both the hadron pair and the muon pair would leave tracks in the detector, this secondary vertex can be reconstructed using the vertexing algorithms discussed in Section 3.2.2.

In this analysis, the Kalman vertex fitter is used which employs the Kalman filter formalism [109–111] that applies a global least-squares minimization to reconstruct the secondary vertex using the track properties of the two selected

charged hadrons or the two selected muons. With the reconstructed secondary vertex, the measured decay length, quantified as the distance of the vertex from the primary vertex, is a useful tool to enhance the sensitivity to long-lived BSM scalar bosons since most of the SM processes would lead to promptly produced tracks.

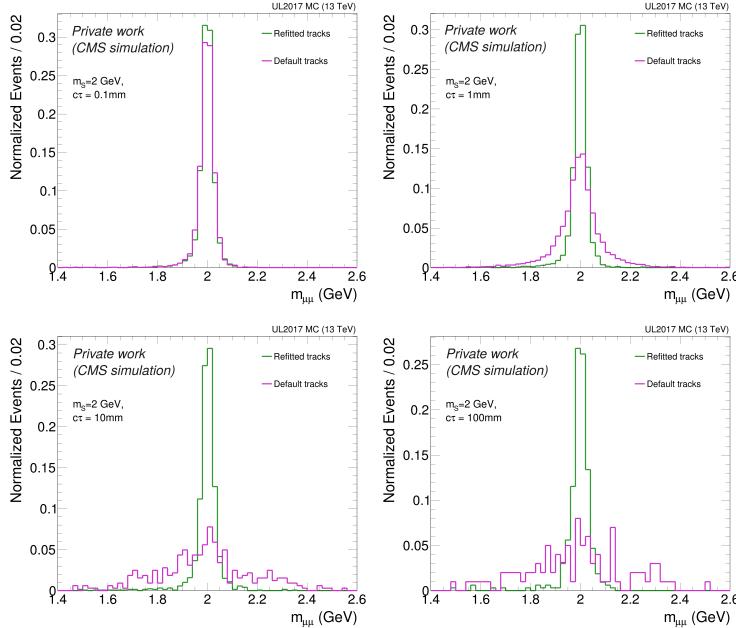


Figure 5.8: Reconstructed dimuon mass distribution of the signal scalars for different lifetimes at $m_S = 2$ GeV using default and refitted tracks.

Once the secondary vertex is reconstructed, the fitter also allows the option to re-evaluate the track properties with secondary vertex as the origin, called smoothing. By default, some tracks originating from long-lived particle decays may be reconstructed using seeds that point towards the primary vertex, which can result in misreconstructed four-momenta. The tracks after the smoothing has been applied are referred to as refitted tracks. Since the refitted tracks more accurately represent the true particle properties at the generator level, a large improvement is seen in the reconstruction of the originating BSM scalar boson. This is shown in Fig. 5.8 where the reconstructed BSM scalar boson mass for the muon final state is shown with the default tracks and the refitted tracks for several

different lifetimes. For the $c\tau = 0.1$ mm, there is no substantial improvement in the dimuon mass resolution since the secondary vertex lies close to the primary vertex. For the larger lifetimes, a significant improvement in the dimuon mass resolution is observed for the refitted tracks, reaching a resolution similar to the $c\tau = 0.1$ mm case. In the analysis, once the secondary vertex is reconstructed, only the refitted tracks for the muons and hadrons are used for further selection criteria.

5.4.2 Analysis Strategy

Following the trigger and building the collection of charged hadrons and muons, the next challenge is the reconstruction of the BSM scalar boson candidates. In the signal samples, the collection of reconstructed particles also contains contributions coming from pileup and the underlying event and ISR which can lead to incorrect BSM scalar boson candidates. For devising the optimal strategy for this, a study with the signal samples is presented here which motivates the baseline selection for the analysis. The signal hypothesis with $m_S = 2$ GeV and $c\tau = 0.1$ mm under 2017 data-taking conditions is used as reference unless specified otherwise. A special sample where both $S \rightarrow \pi^+\pi^-$ and $S \rightarrow K^+K^-$ are included (according to branching ratios from theory [17]) has been used for the study shown in this section.

In Fig. 5.5 (right), it was seen that the final state muons are very collimated with $\Delta R_{\mu\mu} < 0.4$ for most events with at least one muon with high p_T . Since the mass difference between the muon mass ($m_{\mu^\pm} = 0.105$ GeV) and pion or kaon mass ($m_{\pi^\pm} = 0.139$ GeV and $m_{K^\pm} = 0.493$ GeV) is quite small as compared to the momentum of the particles ($p_T > 5$ GeV), similar kinematics are expected for both the muon and charged hadron decay modes. This is shown in Fig. 5.9 and 5.10 where similar kinematics are observed for the muon and charged hadron decay modes. In Fig. 5.10, a longer tail for ΔR_{hh} is present as compared to $\Delta R_{\mu\mu}$ which indicates that due to the mass difference between the hadron and the muon, the hadronic final states are slightly less boosted than the muon final states but the majority of the decays still fall within $\Delta R < 0.4$.

From the generator-level kinematics, there are three major takeaways. The first is that at least one of the muons and one of the hadrons is high p_T , and

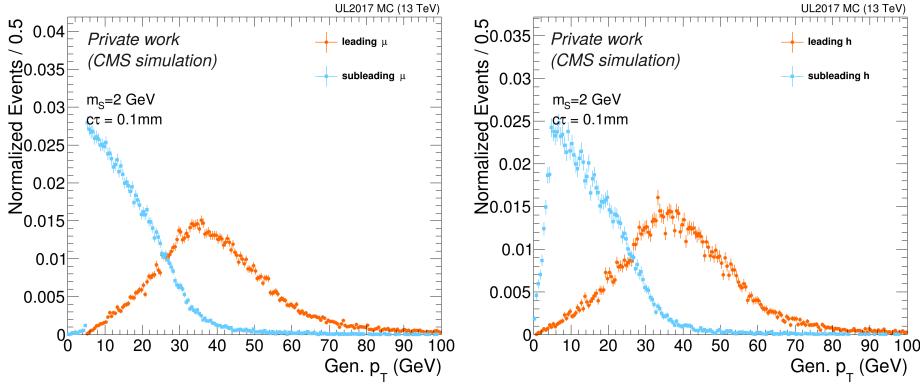


Figure 5.9: Gen. p_T for leading and subleading muon (left) and charged hadrons (right) for $m_S = 2$ GeV, $c\tau = 0.1$ mm.

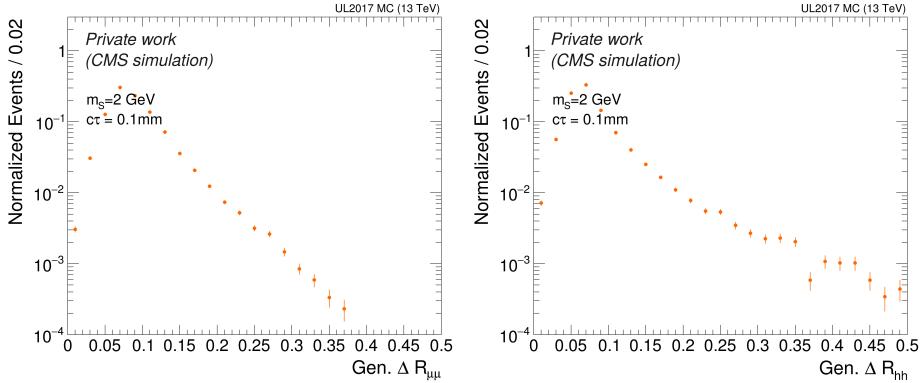


Figure 5.10: Gen. ΔR for $S \rightarrow \mu\mu$ decays (left) and $S \rightarrow hh$ decays (right) for $m_S = 2$ GeV, $c\tau = 0.1$ mm.

since these particles are produced from the hard scattering event, they are likely to be the highest p_T objects in the event. The second is that the angular separation between the two muons or the two hadrons is very small. Finally, the third takeaway is that the kinematics of the hadron and muon final states are similar, implying that a similar selection strategy should be sufficient to reconstruct both the BSM scalar bosons in the event. With this in mind, a preliminary event selection is formulated and listed below.

- Trigger: Isolated single muon trigger (HLT_IsoMu27 for 2017)

- Muons: $n_\mu \geq 2$ from the muon collection (Section 5.4.1) with $p_T > 5$ GeV and $|\eta| < 2.4$. Additionally, $p_T > 29$ GeV on the leading muon is applied to stay in the plateau of the trigger turn-on curve.
- Charged hadrons: $n_h \geq 2$ from the charged hadron collection with $p_T > 5$ GeV and $|\eta| < 2.4$. The charged hadron tracks are assumed to be charged pions.
- Dimuon candidate: The highest p_T -ordered, oppositely charged pair of muons within $\Delta R_{\mu\mu} < 0.4$.
- Dihadron candidate: The highest p_T -ordered, oppositely charged pair of charged hadrons within $\Delta R_{hh} < 0.4$.

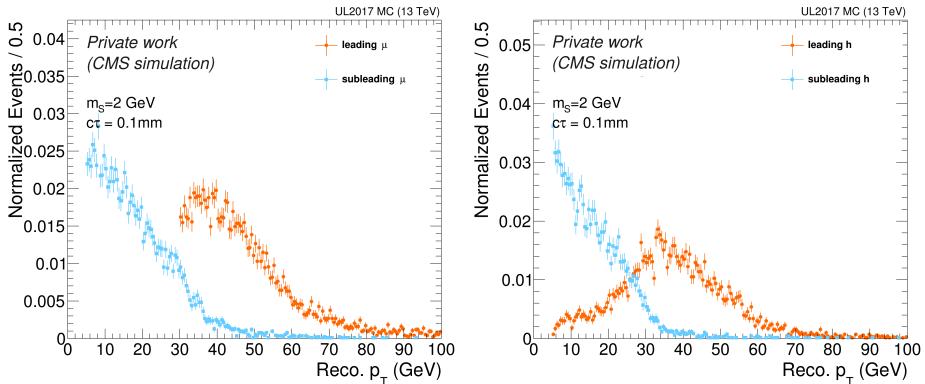


Figure 5.11: Reconstructed p_T for leading and subleading muon (left) and hadron (right) after passing a preliminary event selection.

Following the event selection listed before, the reconstructed p_T for the final state objects (muons and hadrons) is shown in Fig. 5.11. The p_T distributions resemble what is observed at the generator level, with a p_T cut present for the leading muon due to the trigger constraints discussed before. This further on affects the dimuon $\Delta R_{\mu\mu}$ distribution (Fig. 5.12 (left)) which shows that the events passing the selection are more collimated than the events at the generator level since a large p_T threshold implies that the muons are produced with a larger Lorentz boost which in turn causes the muons to be much closer to each other.

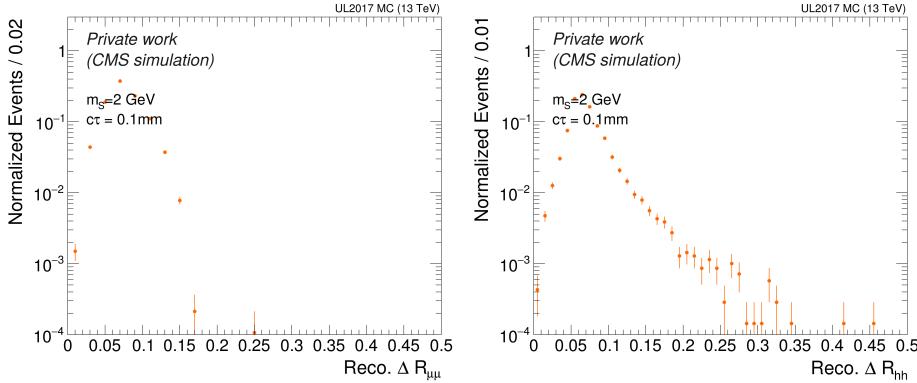


Figure 5.12: Reconstructed ΔR for $S \rightarrow \mu\mu$ and $S \rightarrow hh$ after passing a preliminary event selection.

Comparing the generator level distributions with the reconstructed properties of the final state particles illustrates that the selection is efficient at selecting the targeted particles.

In Fig. 5.13, the reconstructed BSM scalar boson mass using the dimuon (left) and dihadrons (right) are shown. For the dimuon, the peak is seen at $m_S^{\text{reco}} = 2 \text{ GeV}$ as expected with a narrow width. For the dihadron, the distribution is much broader with two prominent peaks at $m_S^{\text{reco}} \sim 1.8$ and 2 GeV . The two mass peaks are expected since the hadronic decay channel in this special signal sample consists of events where the scalar can decay to either $\pi^+\pi^-$ or K^+K^- . Since the default particle assumption for the charged hadrons is pion, the larger peak (with final state K^\pm) is seen to be at a lower m_S while the smaller peak (with final state π^\pm) appears at $m_S^{\text{reco}} = 2 \text{ GeV}$. This is in line with the branching ratio calculations in [17], where the decays to kaons dominates over the decay mode to pions.

Compared to the muon final state, a deeper investigation is required for the hadronic final state since the number of hadrons in the event are much larger than the number of muons (especially with the high p_T cut on the leading muon) which suggests that the possibility of selecting the wrong hadron pair is higher. This is studied in detail by comparing the kinematics of the reconstructed particles for the corresponding generator level hadrons (i.e. gen. matched particles)

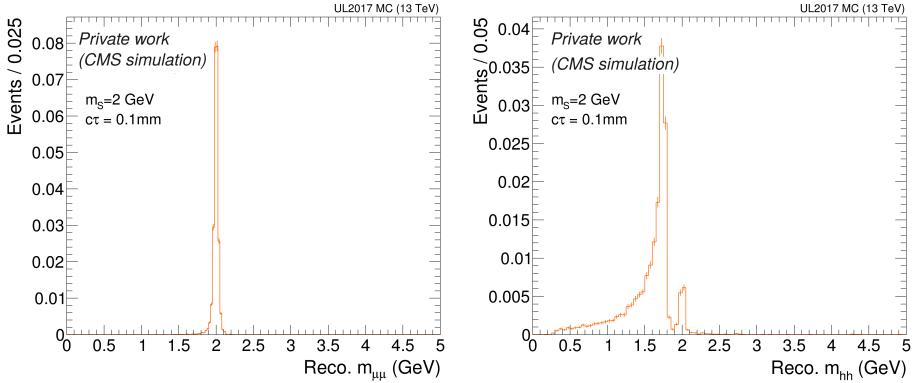


Figure 5.13: Reconstructed scalar mass for $S \rightarrow \mu\mu$ and $S \rightarrow hh$ after passing the preliminary conditions.

with the hadrons selected with the selection procedure outlined before. The gen. matching procedure considers the generator level hadrons produced from the decays of the BSM scalar boson and searches for the closest reconstructed charged hadron track. If a reconstructed hadron is extremely close (defined as $\Delta R < 0.03$ between the gen. particle and the reconstructed particle) then it is considered as the gen. matched particle. The reconstructed p_T for the gen. matched hadrons and the corresponding ΔR distribution between them is shown in Fig. 5.14. Note that the events in Fig. 5.14 are not required to pass the event selection and instead only require the presence of ≥ 2 final state hadrons and ≥ 2 final state muons to provide a collection of all the events that can be reconstructed.

In the next stage, the events passing the preliminary event selection are checked to see how many of them select the gen. matched hadrons. The events where both selected hadrons are matched to the signal hadrons at the generator level (and hence pass the gen. matching procedure) are classified as gen. matched events. The remaining events (failing the gen. matching procedure) are classified as non-gen. matched events. In Fig. 5.15, a comparison between the gen. matched events and non-gen. matched events is shown in the dihadron mass and dihadron p_T distribution. As these non-gen. matched hadrons most likely come from pileup and underlying event, they populate the low values of dihadron p_T and dihadron mass. In the case of the dihadron p_T distribution in Fig. 5.15, it

is seen that the region below 20 GeV is mostly populated with non-gen matched events whereas the gen. matched events are fairly low, indicating that the preliminary selection is not efficient at capturing the signal hadrons in this region. Therefore, a dihadron p_T cut of $p_T > 20$ GeV is included in the event selection to ensure that events with incorrectly selected hadron pairs are rejected.

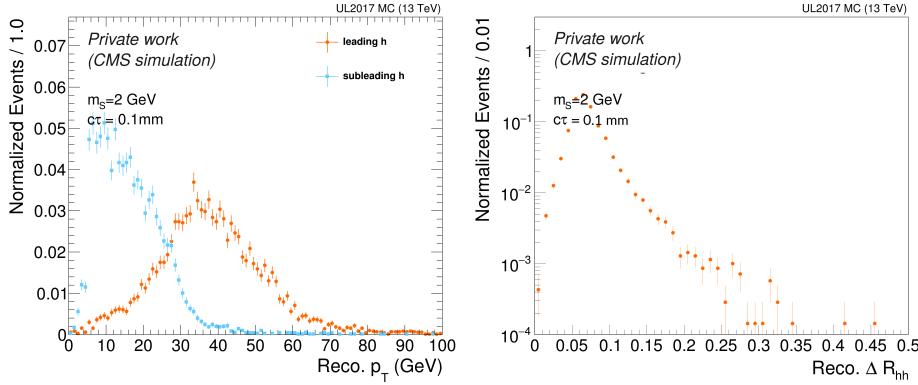


Figure 5.14: Reconstructed p_T for gen. matched leading and subleading hadrons (left) and reconstructed ΔR_{hh} for $S \rightarrow hh$ (gen. matched) (right).

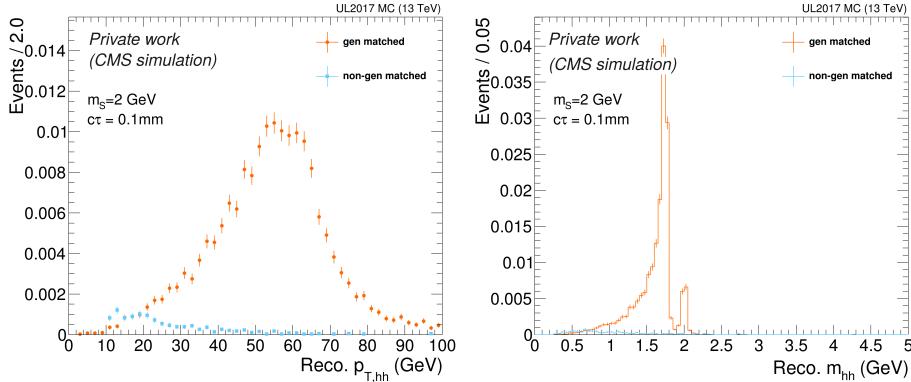


Figure 5.15: Reconstructed BSM scalar boson p_T (left) and mass for $S \rightarrow hh$ (right) using gen. matched and non-gen. matched hadrons for events passing the preliminary selection. Events are normalized to the same total number of events passing the selection.

From Fig. 5.13, it is seen that even if the event selection picks the correct

hadrons (i.e. gen-matched hadrons), the dihadron mass distribution is quite broad compared to the dimuon mass primarily because of the presence of the two peaks coming from the charged kaon and charged pion decay channels. This ends up having a significant impact when the similarity between the reconstructed dihadron and dimuon properties is taken as a handle to reject background events. Additionally, since the gain coming from the charged pion decay mode for $m_S = 2$ GeV is small, the probed hadronic decay mode in the analysis is set to only the dominant one, as discussed in Section 5.2.1. In the case of signal hypotheses with $m_S \geq 1.1$ GeV, the final state consists of two muons and two charged kaons. For $m_S < 1.1$, the final state consists of two muons and two charged pions. This is summarized in the Eqns. 5.5,5.6. The analysis is hence split into two different final states:

$$gg \rightarrow H \rightarrow S\bar{S} \rightarrow \mu^+ \mu^- + K^+ K^- \quad m_S \in [1.1, 2] \text{ GeV} \quad (5.5)$$

$$gg \rightarrow H \rightarrow S\bar{S} \rightarrow \mu^+ \mu^- + \pi^+ \pi^- \quad m_S \in [0.4, 1.1] \text{ GeV} \quad (5.6)$$

For the signal hypothesis with $m_S \geq 1.1$ GeV, the charged hadrons four-momenta are redefined with the charged kaon assumption. Including the requirement of a common vertex and using the subsequent refitted tracks further im-

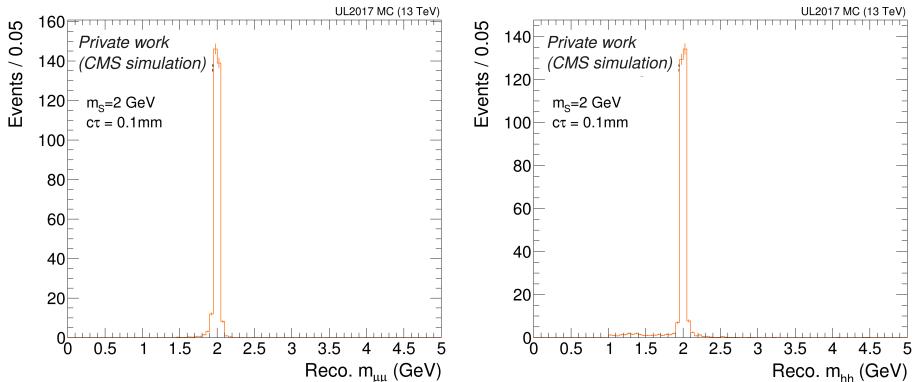


Figure 5.16: Reconstructed scalar mass $S \rightarrow hh$ after applying charged kaon mass assumption for the hadrons and using refitted tracks after preliminary selection.

proves the mass resolution. In Fig. 5.16, the reconstructed dimuon and dihadron mass after passing the preliminary selection and the charged kaon assumption for the hadrons demonstrates that the mass resolution is similar between both muon and kaon pairs. This is an important signal feature which is exploited in the full event selection to improve the signal sensitivity.

5.4.3 Baseline selection

Based on the studies presented in the previous section, a baseline selection is formulated. The complete baseline selection is shown in Fig. 5.17. After the selection of the muon & hadron objects, the first step is the selection of the signal-like objects. Since the selection strategy is similar between the muons and the charged hadrons, the following is applied to both objects.

The selection procedure proceeds with iterating over the p_T -ordered object collection until some conditions are satisfied. The first condition is the presence of a pair with oppositely charged objects with $\Delta R < 0.4$. The next step checks if the objects come from a valid secondary vertex, defined as a common vertex where the fit converges and yields a non-zero χ^2 . If this condition is satisfied,

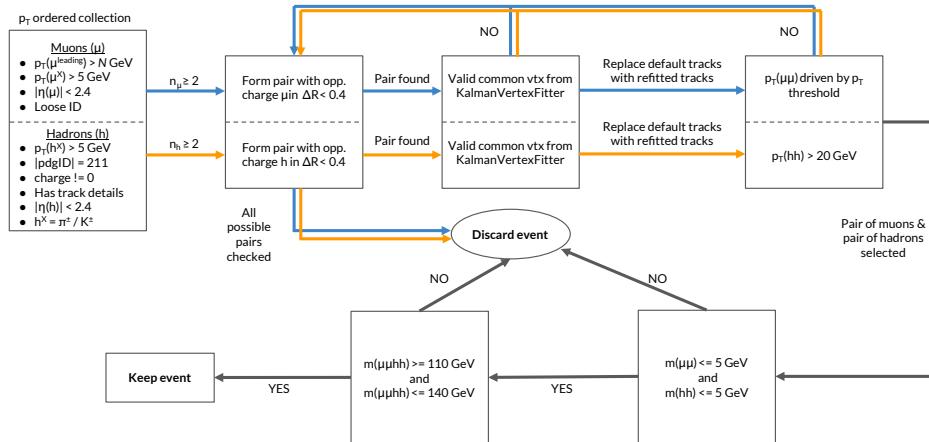


Figure 5.17: Flowchart illustrating the object selection and the baseline selection. The p_T thresholds on the muon collection are modified according to the event trigger as described in Section 5.3. All other cuts across the different eras of Run 2 remain the same.

the default tracks for the objects are replaced with refitted tracks. In the charged hadron channel, an explicit cut of $p_T > 20$ GeV is imposed on the dihadron p_T . For the muon channel, no explicit cut is imposed since the trigger requirements already enforce an implicit dimuon p_T cut. If all these conditions are passed, the iteration over the muon and the charged hadron collection stops. If no dihadron or dimuon object is found, the event is discarded.

Once a pair of muons and a pair of hadrons is selected, their respective vertices are required to pass through additional quality requirements which helps in rejecting events with unphysical attributes such as the uncertainty on the vertex position. Specifically, the covariance matrix of the vertex is used to derive the uncertainty along the direction of the vector connecting the secondary vertex and the primary vertex, in the transverse plane. If this uncertainty is negative, the event is rejected. This cut is also effective in rejecting events coming from background processes.

The next few selection criteria include some loose cuts. The first removes events at large dimuon and dihadron masses by requiring the mass to be less than 5 GeV for both. Using the muons and the charged hadrons, the Higgs boson candidate is reconstructed. The four-object invariant mass, $m_{\mu\mu hh}$, is required to be near the SM Higgs boson mass by applying the condition $110 < m_{\mu\mu hh} < 140$ GeV to suppress the backgrounds from low-energy processes.

5.4.4 Signal Region & Control Region

So far, the event selection criteria have been guided by the signal simulation and include only loose cuts, which are chosen to reduce the background without any impact on the signal acceptance. In the following sections, additional selections are introduced that have been optimized to keep a good signal acceptance while significantly reducing the background. To design these selections, the relevant background processes are studied using a simulation-based approach and a data-driven approach. For the simulation-based approach, the Standard Model background processes listed in Table 5.5 are used. For the data-driven approach, special care needs to be taken to ensure that the data being used for this do not contain potential signal events to avoid any bias in tuning the selection. This is achieved by masking or blinding the region where the signal is expected (the

signal region or SR). Conversely, a control region (CR) is defined to be rich in background while containing events with similar kinematics as the signal events. This can then be used to study the background contribution directly from data.

For the definitions of the signal region and the control region, the reconstructed four-object mass distribution, $m_{\mu\mu\text{hh}}$, is used. In the case of the signal, $m_{\mu\mu\text{hh}}$ represents the reconstructed Higgs boson mass, and the distribution is expected to be concentrated near $m_{\mu\mu\text{hh}} = 125$ GeV. A window of ± 2.5 GeV around the Higgs boson mass is taken as the signal region (i.e. $m_{\mu\mu\text{hh}} \in [122.5, 127.5]$ GeV). Fig. 5.18 shows the $m_{\mu\mu\text{hh}}$ distribution for the 2017 data, SM simulations and two reference signal samples, under the charged pion assumption (right) and the charged kaon assumption for the final state hadrons (left). For a direct comparison with data, the simulations are scaled to the luminosity of the 2017 dataset $\mathcal{L}_{\text{int}} \sim 41.5 \text{ fb}^{-1}$. Additionally, the simulation samples are corrected with relevant data-to-simulation corrections to account for discrepancies with respect to the data. A discussion of the corrections and their associated uncertainties is carried out in Section 5.7. As a result, a reasonably good agreement is observed between the data and simulation in the sidebands of the $m_{\mu\mu\text{hh}}$ distribution.

In the signal region, the dominant background consists of multijet events from quantum chromodynamic (QCD) processes, followed by some contribution from $t\bar{t}$ production. The QCD multijet background is dominated by events with heavy-flavour jets, which can yield muon pairs in the jet hadronization. Along with the data in the sidebands of the $m_{\mu\mu\text{hh}}$ distribution, the backgrounds follow an exponentially falling nature. This is denoted by the fit on data which is extrapolated to the signal region and shows a good compatibility with the SM simulations.

This leads to the definition of the control region for data. The control region is chosen to be the sidebands of the $m_{\mu\mu\text{hh}}$ distribution and is defined as $m_{\mu\mu\text{hh}} \in [110, 122.5] \cup [127.5, 140]$. Since it is seen that the backgrounds follow an exponentially decaying nature in the $m_{\mu\mu\text{hh}}$ distribution, a fit on the sidebands allows for the extrapolation to the signal region and provides an initial estimate of the expected background in the signal region. A transfer factor, $f = N_{\text{fit}}^{\text{SR}} / N_{\text{fit}}^{\text{CR}}$, can be applied to the events in the control region to provide the rough back-

ground distribution of other kinematic variables. Note that this is not the final background estimation; This is used to provide an indication of the signal-to-background discrimination during different stages of the selection. The final background yield is calculated from the procedure discussed in Section 5.5 after the full event selection is applied. Additionally, during the formulation of the event selection, a tight control region of $m_{\mu\mu hh} \in [120, 122.5] \cup [127.5, 130]$ GeV is chosen to have a more accurate description of the kinematics of background in the signal region. Another benefit of this is that the transfer factor is approximately 1 while the transfer factor for the looser control region is ~ 0.2 .

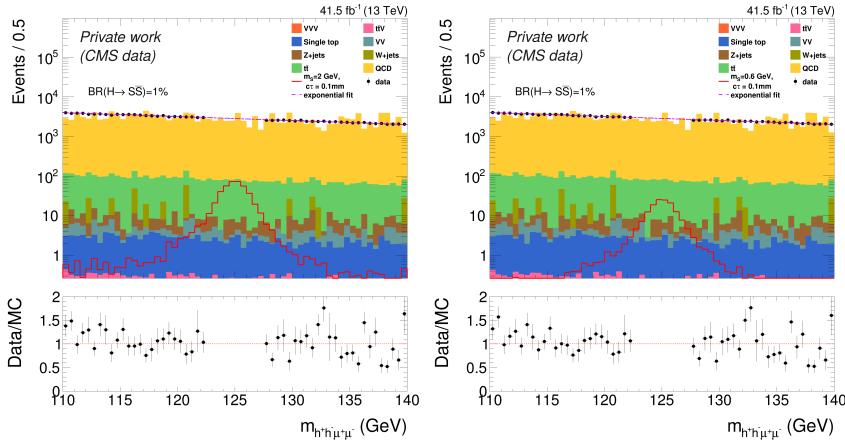


Figure 5.18: Reconstructed four-object mass distribution ($m_{\mu\mu hh}$) with blinded data, lumi-scaled SM samples and signal samples with $h^\pm = K^\pm$ (left) and $h^\pm = \pi^\pm$ (right). The signal samples for $m_S = 2$ GeV, $c\tau = 0.1$ mm and $m_S = 0.6$ GeV, $c\tau = 0.1$ mm are shown as reference. An exponential fit($= Ne^{\alpha m_{\mu\mu hh}}$) is performed on the Higgs sideband in data to derive the transfer factors from the loose and tight control regions.

Overall, it is seen that the signal region is submerged with background processes, indicating that more selection cuts are required to find any potential signal events. These additional selection criteria are discussed in the following sections. In all the following plots, the background processes (simulation and data) are shown in the tight control region while the signal is shown in the signal region, according to the definitions given above.

5.4.5 Reconstructed scalar mass consistency

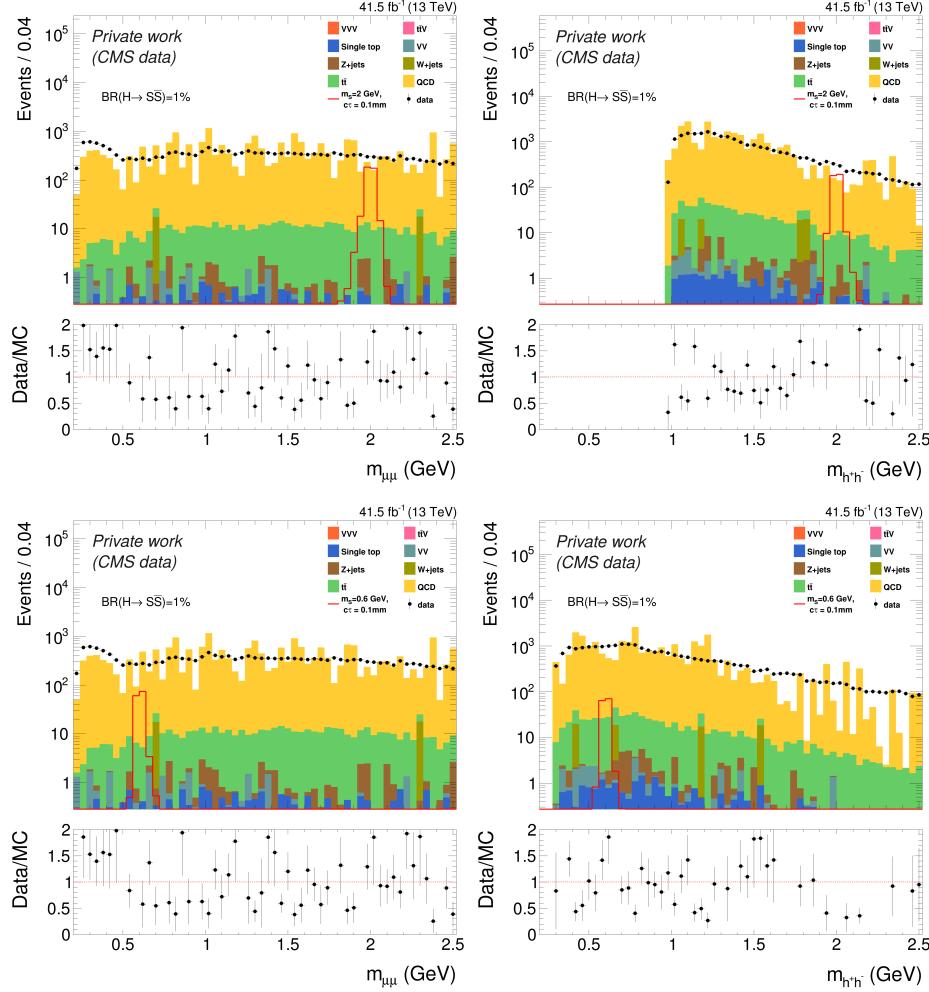


Figure 5.19: Reconstructed dimuon (left) and dihadron (right) mass distribution with 2017 data, background simulation in the tight control region (i.e. $m_{\mu\mu hh} \in [120, 122.5] \cup [127.5, 130]$ GeV) and signal simulation ($m_{\mu\mu hh} \in [122.5, 127.5]$ GeV) for the assumptions, $h^\pm = K^\pm$ (top row) and $h^\pm = \pi^\pm$ (bottom row). The signal samples for $m_S = 2$ GeV, $c\tau = 0.1$ mm and $m_S = 0.6$ GeV, $c\tau = 0.1$ mm respectively, are shown as reference.

In Fig. 5.18, it is seen that a similar number of background events are ob-

served for the charged pion assumption (i.e. $h^\pm = \pi^\pm$) and the charged kaon assumption (i.e. $h^\pm = K^\pm$) cases since the redefinition of the charged hadron does not impact the background rejection efficiency. This is because the mass difference between the kaon and the pion is small compared to the p_T of these final state particles. However, a significant difference is observed when considering the dihadron candidate mass in both cases. This is shown in Fig. 5.19 where the reconstructed dimuon and dihadron mass is shown with the charged kaon assumption (top) and the charged pion assumption (bottom) after the baseline selection and blinding is applied on the 2017 datasets. For the dihadron mass distribution, the background increases towards low masses since it consists of low- p_T hadrons. The edge of the dihadron invariant mass distribution however shifts to $m_{hh} \sim 1$ GeV for the charged kaon assumption due to the constraint on the minimum possible dikaon invariant mass ($m_{\pi^\pm}, m_{K^\pm} = 0.139, 0.493$ GeV). For the dimuon mass distribution, the background is also observed to increase towards low masses with a broad resonance around ~ 0.3 GeV which comes from $\phi \rightarrow K^+ K^-$ events where the kaons are misidentified as muons [151].

While the background is spread in the entire dimuon and dihadron mass range considered in Fig. 5.19, the signal events are localized within narrow peaks centered on the corresponding signal mass hypothesis, with similar resolution in both the muon and hadron channels. This similarity between the reconstructed dihadron mass and dimuon mass provides a powerful handle to reduce the background. In this analysis, a selection is defined in the two-dimensional (2D) plane of the dimuon and dihadron masses which correlates the two quantities.

The first step in this method is to extract the mass resolution of the reconstructed dimuon and dihadron objects for all the BSM scalar boson mass points considered. Since it was seen in Fig. 5.8 that the secondary vertex reconstruction leads to similar mass resolution between the different lifetimes, only the prompt signal samples have been chosen for extracting the resolution. In Fig. 5.20, a Gaussian fit is shown on the dimuon and dihadron mass distributions for $m_S = 2$ GeV. The spectrum is distorted from a perfect Gaussian function due to possible photon radiations, material interactions and reconstruction failures which lead to non-Gaussian tails. As a result, the Gaussian has been fit only to the core of the distribution. After the fit, the dimuon and dihadron mass resolutions

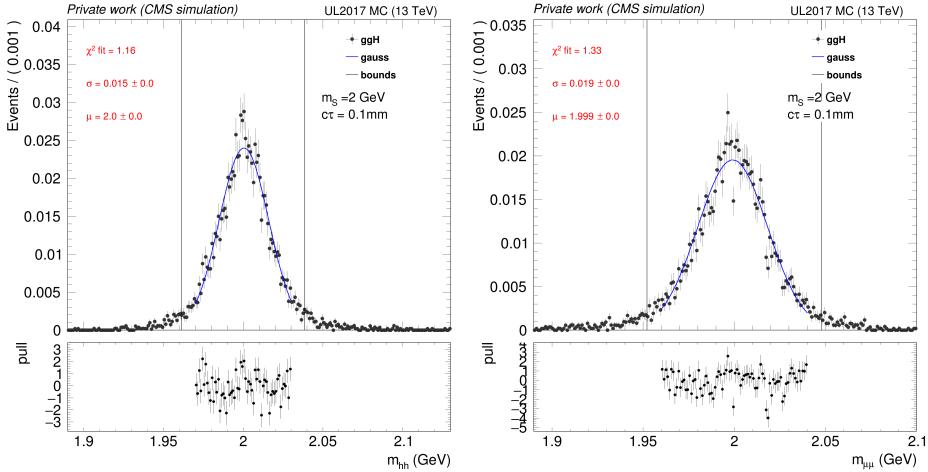


Figure 5.20: Dimuon (left) and dihadron (right) mass distributions for $m_S = 2$ GeV, $c\tau = 0.1$ mm with a Gaussian fit on the core of the distribution. The corresponding mass bounds selected by choosing $\mu \pm 2.5\sigma$ from the fit are shown as black lines.

(σ_{fit}) have been used to define mass bounds on the dimuon and dihadron mass by taking $\pm 2.5\sigma_{\text{fit}}$ around the mean. This is expected to contain most of the signal events, as seen in Fig. 5.20. This is then repeated for each of the mass points considered in the analysis with the appropriate hadron mass assumption ($h^\pm = K^\pm$ for $m_S \geq 1.1 \text{ GeV}$ and $h^\pm = \pi^\pm$ for $m_S < 1.1 \text{ GeV}$). With these pairs of dimuon and dihadron mass bounds, a bounding box in the 2D plane of dihadron and dimuon mass is constructed for each mass point, as seen in Fig. 5.21. A linear extrapolation is performed on the outer edges of the bounding boxes to produce a mass selection in this 2D plane where the events falling within the lines are accepted. As the scalar mass gets smaller, the mass selection becomes narrower. Fig. 5.22 shows the robustness of the mass window with respect to the lifetime for $m_S = 2$ GeV. As a result of the secondary vertex reconstruction and refitted tracks, the samples with the largest displacement are still well-reconstructed with a resolution similar to what is present for the prompt samples.

To understand the impact of the selection on the background, Fig. 5.23 shows the 2D dimuon and dihadron mass distribution for QCD simulation using the

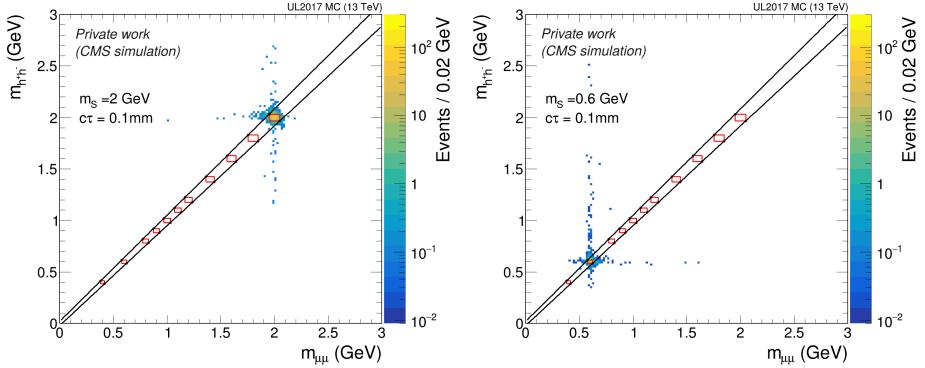


Figure 5.21: Two-dimensional plane of the reconstructed dihadron mass versus dimuon muon for a BSM scalar boson of $m_S = 2 \text{ GeV}$ (left) and $m_S = 0.6 \text{ GeV}$ (right) with proper lifetime $c\tau = 0.1 \text{ mm}$. The red boxes denote the bounding boxes formed by taking the bounds after the Gaussian fit for each of BSM scalar boson hypothesis. The black lines denote the mass bounds as a function of the dimuon and dihadron mass, which would accept the events falling inside it.

kaon mass assumption and the pion mass assumption for the final state hadrons. The sharp cut-offs at $m_{\mu\mu} \sim 0.2 \text{ GeV}$ and $m_{hh} \sim 1 \text{ GeV}$ in Fig 5.23 (left) correspond to the minimally possible dimuon and dikaon masses. Similarly, for Fig. 5.23 (right), the cut-off in m_{hh} is driven by the minimum possible dipion mass.

At this stage, it is also useful to note that the QCD simulation samples are already statistically limited, as seen by the sporadic distribution of the events. This will further reduce with the addition of more selection criteria and result in a small number of events with large weights. This already motivates that a data-driven approach to estimating the background is required since the SM simulation samples would come with large uncertainties.

In Fig. 5.24, the 2D mass distribution is shown for the 2017 data in the tight control region (i.e. $m_{\mu\mu hh} \in [120, 122.5] \cup [127.5, 130] \text{ GeV}$). The background is concentrated at low dihadron masses and is uncorrelated with the dimuon mass. The application of the 2D mass selection leads to a reduction of the background by $\sim 96\%$, therefore providing strong signal-background discrimination.

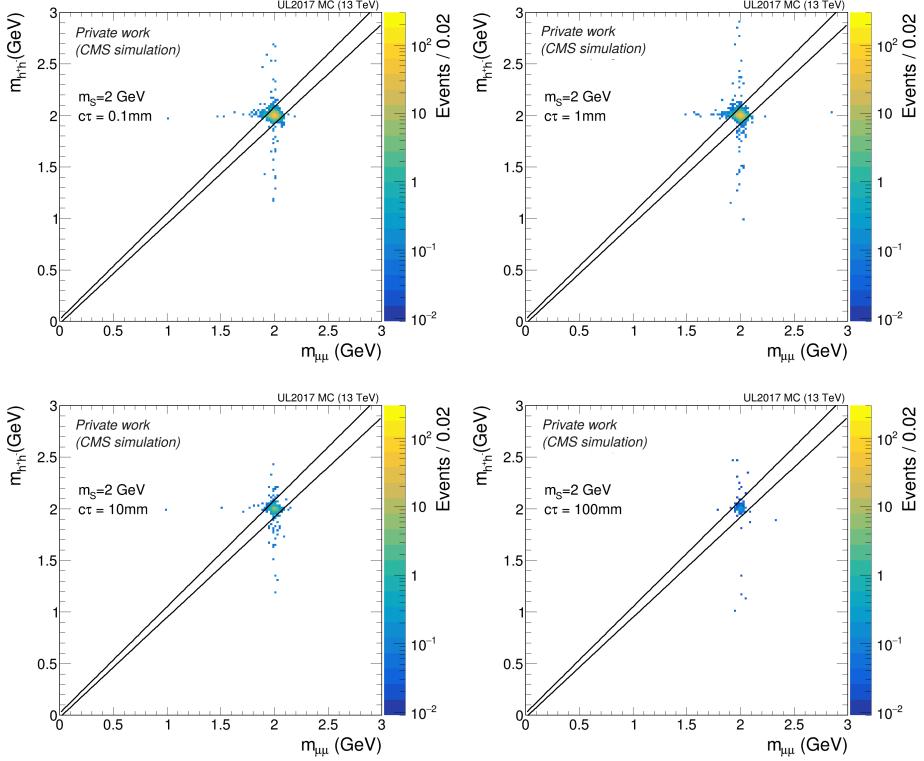


Figure 5.22: Performance of the mass bounds in accepting signal samples with $c\tau = 0.1$ (top left), 1 (top right), 10 (bottom left), 100 (bottom right) mm for $m_S = 2$ GeV.

With the application of the 2D diagonal mass cut defined here, the reconstructed dimuon and dihadrons masses are ensured to be similar, in both signal and the background processes. Instead of studying the the dimuon and dihadron mass distributions separately as the reconstructed masses of the BSM scalar boson candidates, the average di-object mass, $\langle m_{\mu\mu, hh} \rangle$, is a robust estimate of the BSM scalar boson mass at the reconstructed level, after the selection is applied. This quantity is useful for extracting the background and signal yield and is discussed further in Section 5.5.

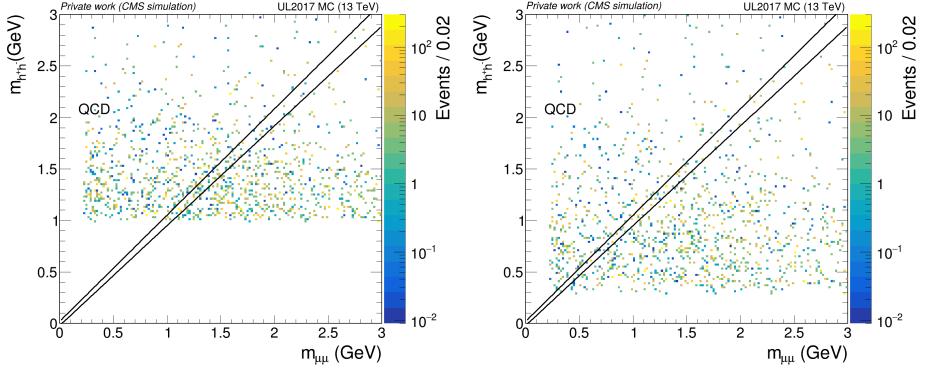


Figure 5.23: Two-dimensional plane of the reconstructed dihadron mass versus dimuon muon for QCD simulation in tight CR with charged kaon assumption (left) and charged pion assumption (right) for the final state hadrons. Sporadic distribution of events, visible due to limited statistics of the samples.

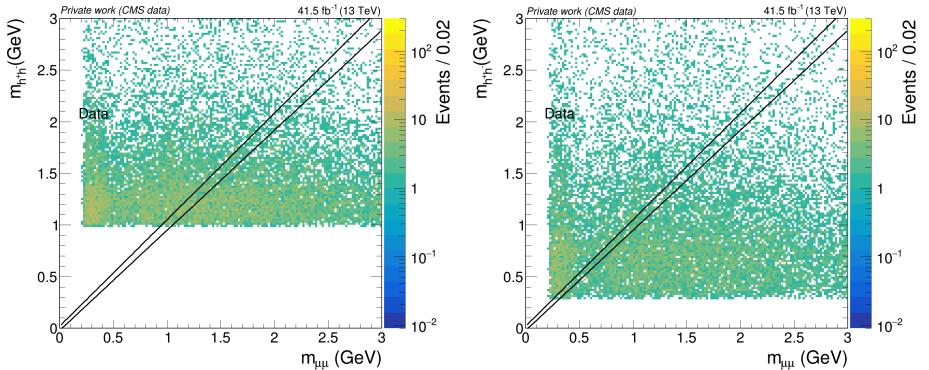


Figure 5.24: Two-dimensional plane of the reconstructed dihadron mass versus dimuon muon for the 2017 data in tight CR with charged kaon assumption (left) and charged pion assumption (right) for the final state hadrons.

5.4.6 Transverse displacement and Event categories

During the baseline selection, the secondary vertices (SV) are reconstructed from which the final state muons and charged hadrons originate. Besides providing access to improvement in the track 4-momenta description (through track refitting), their position can be used to boost the sensitivity to BSM scalar particles with

larger lifetimes. Signal samples with $c\tau \geq 1$ mm are expected to have a substantial number of events with distinct secondary vertices which lie away from the primary vertex. Since the dominant background in the analysis comes from multijet QCD processes which mostly contain particles produced at the primary vertex (PV), a measure of the distance from the primary vertex for the secondary vertex can provide a good handle to optimize the signal to background ratio. In particular, by introducing event categories based on the information from the secondary vertices allows for separation of background rich region from the signal rich regions. This is further discussed in the section below, starting with the definition of the measure to draw the event categories.

The Kalman vertex fitter provides the secondary vertex position as well as a 3×3 covariance matrix of the secondary vertex fit. Since the Higgs boson is predominantly produced centrally, it is expected that most of the boost is present in the transverse plane. Three quantities, based on the position in the transverse plane, are therefore used in the analysis:

- L_{xy} – Transverse distance travelled by the reconstructed dimuon or dihadron. This is measured as the distance between the secondary vertex and primary vertex in the transverse plane.
- σ_{xy} – Total uncertainty on the primary vertex and secondary vertex reconstruction in the transverse plane. This is calculated from the covariance matrix of the SV and PV in the direction of the PV-SV axis.
- $\frac{L_{xy}}{\sigma_{xy}}$ – Termed as L_{xy} significance. This is a measure of the transverse displacement relative to the uncertainty on the measurement and takes into account the detector resolution.

In the baseline selection, a cut on the σ_{xy} was included to remove events with unphysical vertices. Such unphysical vertices can occur when the tracks are very close to each other but do not originate from a common vertex. In such a scenario, the fitting may converge but provide a vertex position lying far beyond the detector volume. One of the indicators for this is its effect on the covariance matrix. While the matrix should be symmetric, in some cases it is found to be asymmetric. In such cases, the covariance matrix is made symmetric by comparing the off-diagonal elements and assigning the value of the element which

results in a larger overall variance (i.e. σ_{xy}^2) to ensure a conservative estimate. There are still however events where the matrix leads to a negative variance, in which case, the event is discarded by the baseline selection cuts.

In the analysis, two secondary vertices are reconstructed – the dihadron vertex and the dimuon vertex. Figure 5.25 shows the L_{xy}^{hh} (left) and the $L_{xy}^{\mu\mu}$ (right) for the charged kaon assumption for the final state hadrons with the reference signal sample $m_S = 2$ GeV and $c\tau = 10$ mm, 2017 data and the SM background simulation after the baseline selection. A sharp cutoff for the signal events is observed for $L_{xy}^{\mu\mu}$ due to the d_{xy} cut on the leading muon applied by the isolated single muon trigger. For both $L_{xy}^{\mu\mu}$ and L_{xy}^{hh} , it is seen that there are events in the background that can extend to very large values of transverse displacement, most likely due to a random pair of tracks that cross each other and give the illusion of a real secondary vertex.

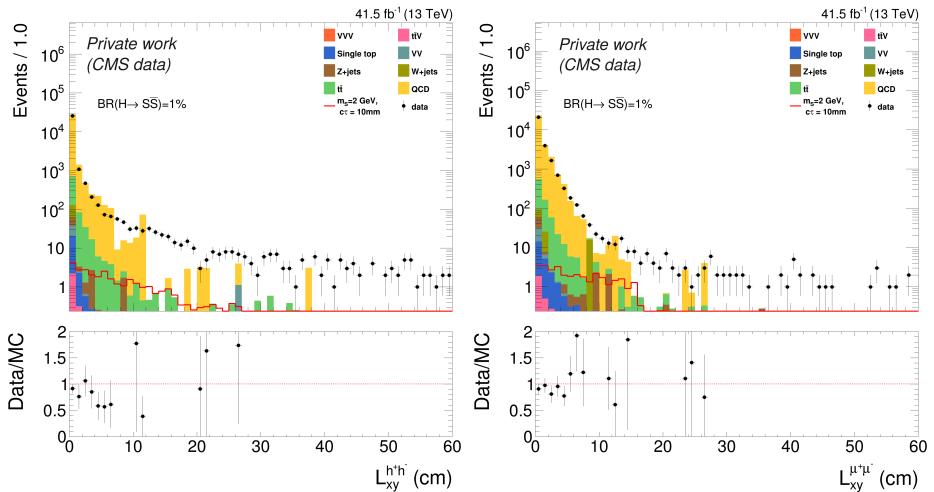


Figure 5.25: L_{xy} for the selected dihadron (left) and dimuon (right) after applying baseline selection with $h^\pm = K^\pm$. The signal sample for $m_S = 2$ GeV, $c\tau = 10$ mm is shown as reference.

This is also present in the $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}}$ and $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}}$ distribution for the same datasets, as shown in Fig. 5.26. For the signal sample, it is seen that the $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}}$ distribution shows a smaller tail than the $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}}$ distribution which is present due to a better SV

reconstruction (and a smaller error) for the dimuon as compared to the dihadron. Since the features of both L_{xy} and σ_{xy} are present in L_{xy} significance, $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}}$ and $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}}$ are used to quantify the distance travelled by the reconstructed dimuon and dihadron objects. As these quantities are correlated and depend on the lifetime of the BSM scalar boson for the signal, the 2D plane of $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}}$ and $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}}$ is used to study the features of the signal and background events.

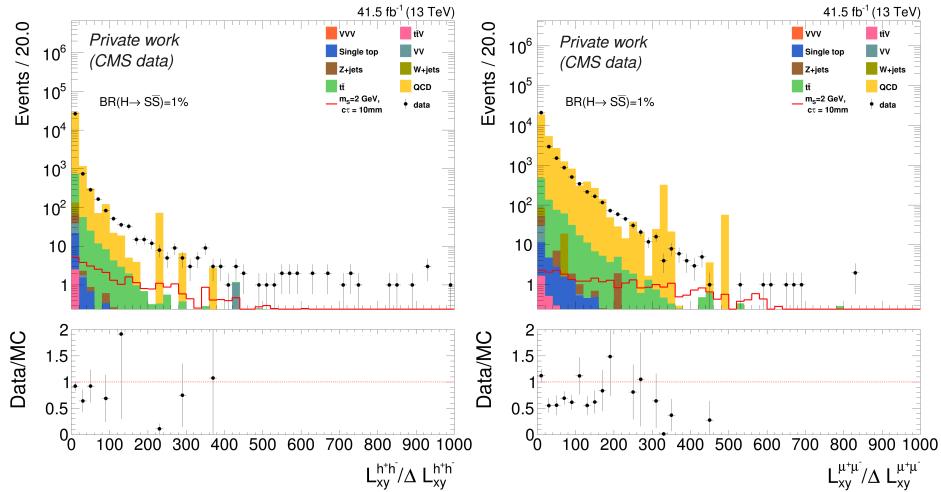


Figure 5.26: L_{xy} significance for the selected dihadron (left) and dimuon (right) after applying baseline selection with $h^+ = K^+$. The signal sample for $m_S = 2$ GeV, $c\tau = 10$ mm in SR is shown, along with 2017 data and background simulations in tight CR.

Fig. 5.27 shows the 2D distribution of $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}}$ and $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}}$ for signal samples with $m_S = 2$ GeV and $c\tau = 0.1, 1, 10, 100$ mm. As the lifetime increases, the distribution of events spreads out in this 2D plane, extending along each axis. This is shown for the simulation with 2017 data-taking conditions. Figure 5.28 gives a closer look at this 2D distribution for the signal with $m_S = 2$ GeV and $c\tau = 1$ mm (left) in comparison to the expected background using the 2017 data in the tight CR (right). It is observed that the background is expected to be constrained mostly to a region where both the dimuon and the dihadron are not displaced. This is followed by background contribution that lies along the $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}}$ or the $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}}$ axis in Fig. 5.28 but not both. A sparsely populated region is present when both

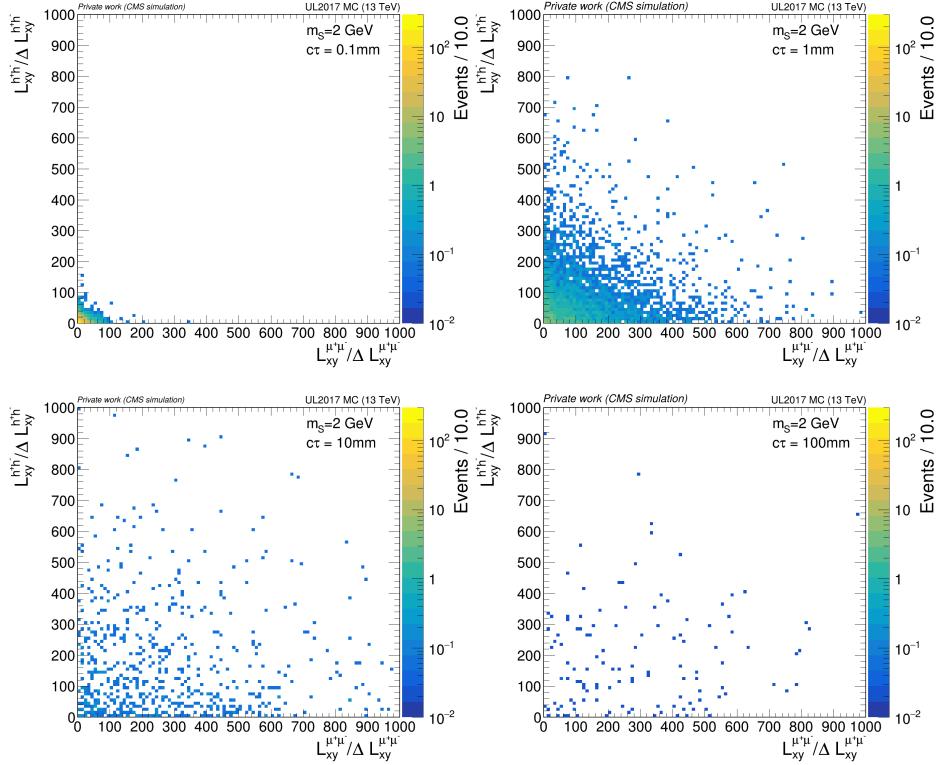


Figure 5.27: 2D distributions of L_{xy}^{hh} significance vs $L_{xy}^{\mu\mu}$ significance for signal samples $m_S = 2$ GeV and $c\tau = 0.1$ (top left), 1 (top right), 10 (bottom left), 100 (bottom right) mm.

the vertices are very displaced, which greatly increases the sensitivity to signal models with $c\tau \geq 1$ mm. In this context, the events are split into 4 categories shown by bounded lines in Fig. 5.28 and has been explained below:

- *prompt*: $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}} < 40$, $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}} < 30$; This category targets to contain $> 90\%$ of the events for the signal benchmarks with $c\tau < 0.1$ mm. The asymmetry in the boundaries is chosen as a result of the difference between the distribution of the two quantities. This is also the region with the most amount of background ($\approx 80\%$ of the background (Fig. 5.28)) and would require tight cuts to improve signal-background discrimination.
- *displaced mu mu*: $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}} > 40$, $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}} < 30$; This category captures events where

the dimuon vertex is displaced while the dihadron vertex is prompt-like.

- *displaced hh*: $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}} < 40$, $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}} > 30$; This category includes events where the dihadron vertex is displaced while the dimuon vertex is prompt-like.
- *displaced*: $\frac{L_{xy}^{\mu\mu}}{\sigma_{xy}^{\mu\mu}} > 40$, $\frac{L_{xy}^{hh}}{\sigma_{xy}^{hh}} > 30$; This category contains events where both dihadron and dimuon vertices are displaced and contributes significantly to the sensitivity for signal benchmarks with $c\tau \geq 1$ mm. The least amount of background is present in this region.

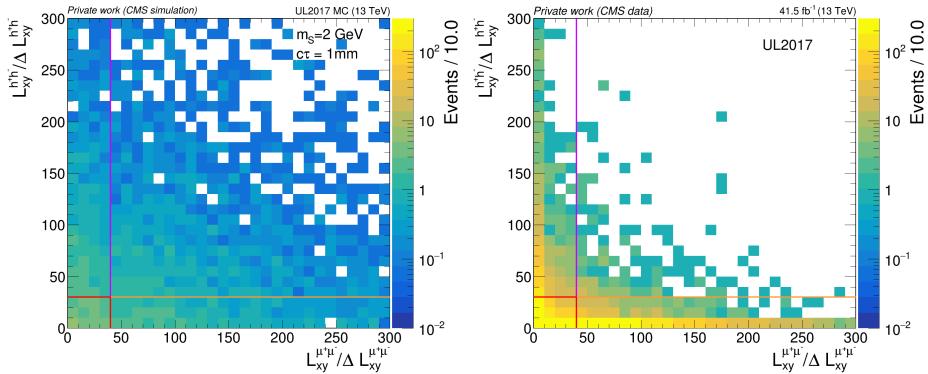


Figure 5.28: 2D distribution of $L_{xy}^{hh}/\Delta L_{xy}^{hh}$ significance vs $L_{xy}^{\mu\mu}/\Delta L_{xy}^{\mu\mu}$ significance for signal sample $m_S = 2$ GeV, $c\tau = 1$ mm and the 2017 data in tight CR. The *prompt* category is denoted with the box enclosed in red lines. The other categories follow once the *prompt* category is made and are shown with the other lines.

5.4.7 Relative Isolation for muons & hadrons

A final selection handle is placed on the amount of hadronic activity around the selected muons and hadrons. One of the ways to quantify this is to measure the isolation around each object. Since the muons and charged hadrons are the only final state particles coming from the hard-scattering process, it is expected that the remaining activity in the event consists of low p_T particles. In particular, since pp collisions are dominated by hadronic activity, requiring the selected muons and hadrons to have small isolation values restricts background events arising from objects like jets.

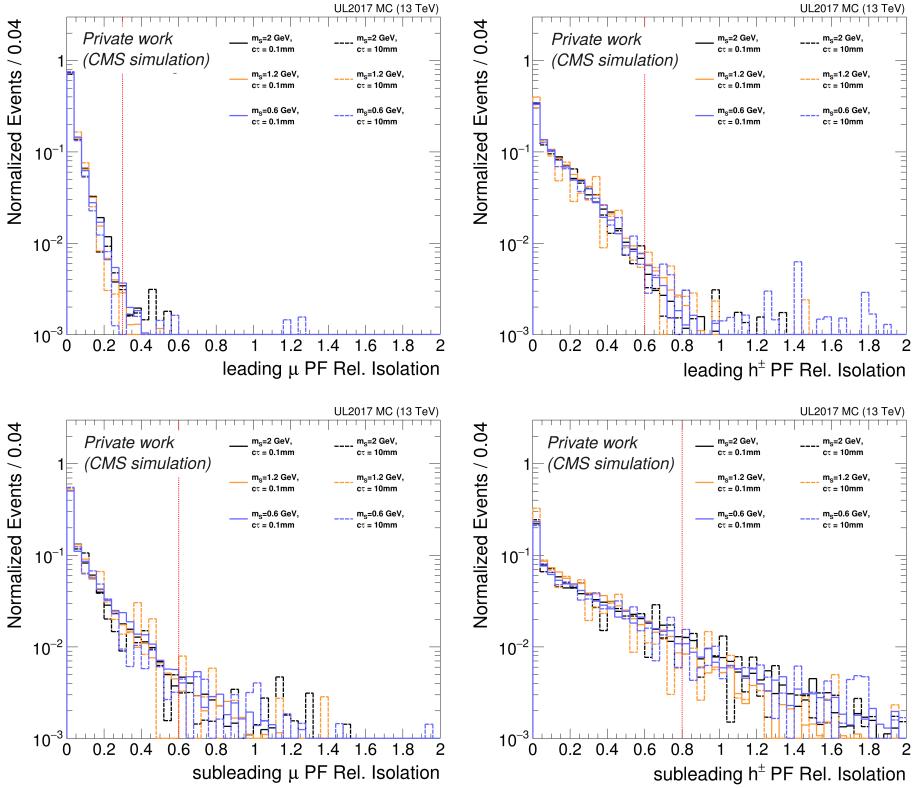


Figure 5.29: PF relative isolation for leading (top) and subleading (bottom) muons (left) and hadrons (right) for $m_S = 0.6, 1.2, 2$ GeV, and $c\tau = 0.1, 10$ mm after passing the baseline selection is shown. The vertical lines denote the isolation cuts applied in the *prompt* category as quoted in table 5.7.

In this context, the PF relative isolation was briefly described in Section 3.2.6. While it is commonly used for muons using Eqn. 3.1, the same can be measured for the charged hadrons with some modifications. Taking the leading charged hadron as an example, the PF relative isolation, according to Eqn. 3.1, would consider the subleading charged hadron as a part of the unwanted hadronic activity around the leading hadron, therefore biasing the calculated isolation. Accordingly, during the calculation of the PF relative isolation for the leading hadron, it is ensured that the subleading hadron is not included in the calculation. Similarly, for the subleading hadron PF relative isolation, the leading hadron p_T is

left out.

Following the modification above, the PF relative isolation is calculated for the leading muon, subleading muon, leading charged hadron and the subleading charged hadron. The relative isolation distributions for the muons (left) and the hadrons (right) are shown in Fig 5.29 for various signal samples. The PF relative isolation peaks at zero for all objects with leading objects being limited to smaller values of relative isolation as compared to the subleading ones. Between the muons and hadrons, the muons are observed to be more isolated as compared to the hadrons, even though they share the same kinematic features. There are several possible reasons for the discrepancy, primarily pertaining to the larger likelihood of misreconstruction and material interactions for the charged hadrons. Finally, it is also observed that the PF relative isolation follows the same shape irrespective of the hypothetical BSM scalar boson mass. Taking note of the differences between the PF relative isolation distribution for the muons and hadrons, an optimal cut is designed to tackle primarily the large backgrounds in the *prompt* category. These cuts are already shown in Fig 5.29 as red lines and their optimization procedure will be discussed below.

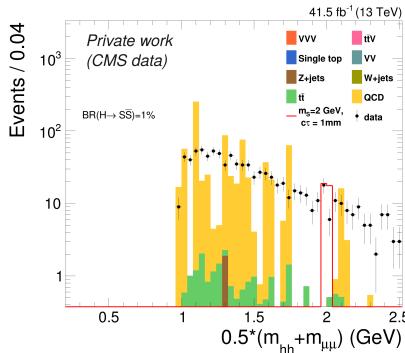


Figure 5.30: Average scalar mass distribution in *prompt* category with signal sample $m_S = 2$ GeV, $c\tau = 1$ mm as reference after applying the 2D diagonal mass cut and the baseline selection. The 2017 data and background simulation from tight CR are also shown.

Since the signal acceptance for the cut does not depend on the BSM scalar boson mass, the optimization for the PF relative isolation cuts is carried out for

the signal sample with $m_S = 1.2$ GeV and $c\tau = 1$ mm. The reason for this choice follows from Fig. 5.30 where the $\langle m_{\mu\mu, hh} \rangle$ distribution is shown for the *prompt* category after the application of the baseline selection and the 2D diagonal mass cut. The background peaks at ~ 1.2 GeV and hence, requires the largest signal-to-background discrimination for suppressing the background.

The strategy to choose the isolation cuts is to optimize for the *prompt* category to reach a ‘tight’ isolation criteria and assigning a ‘loose’ criteria to the remaining displaced categories. Note that the relative isolation for the leading and subleading objects are also correlated since they are nearby each other and hence share similar hadronic activity. Therefore, the leading and subleading object PF relative isolation are studied together in two-dimensional histograms to evaluate the isolation cuts. To evaluate the sensitivity at various values in the two-dimensional PF relative isolation map between the leading and subleading objects, the test statistic $Z = S/\sqrt{B}$ is used as an estimator of the signal significance for the signal. Here, S is the number of expected signal events while B is the expected number of background events as predicted from the data in the tight CR. This test statistic is only applicable when B is large and accounts only for the statistical component of a Poisson distributed background. While this is not a rigorous formulation for performing a statistical interpretation of the discovery potential of the analysis itself, since it does not account for systematic uncertainties, it provides a good base to compare the relative sensitivity at different isolation thresholds.

The selected cuts on the isolation are mentioned in Table 5.7, and Fig. 5.31 shows the isolation distributions in the *prompt* category for the final state muons and hadrons after applying the 2D diagonal mass cut discussed before. While the selected cuts have been chosen based on the distribution for the charged kaon hypothesis, the distribution in principle does not change with the charged pion hypothesis for the selected hadrons since the number of background events still remain the same and the only difference appears in the spread of the events in the reconstructed scalar mass distributions (Fig. 5.24).

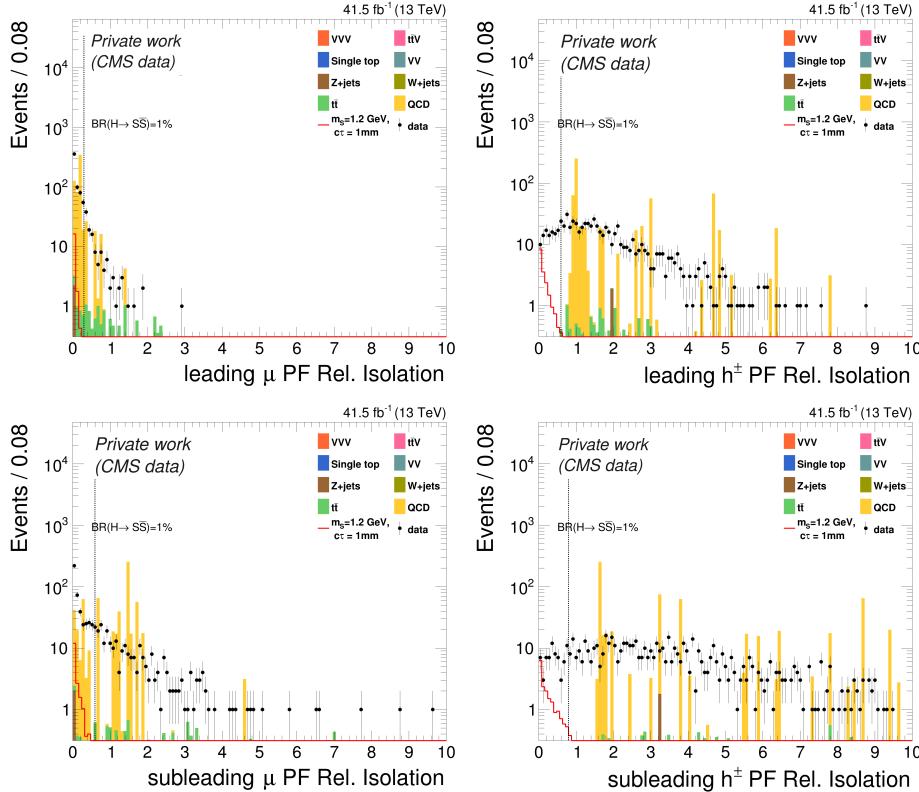


Figure 5.31: PF relative isolation for leading (top) and subleading (bottom) muons (left) and hadrons (right) in the *prompt* category for $m_S = 2$ GeV, $c\tau = 1$ mm after passing the baseline selection and the 2D diagonal mass cut. The 2017 data and background simulation from tight CR are also shown.

5.4.8 Event selection summary

A compact representation of the selection criteria is given in Table 5.8. As an example, Fig. 5.32 shows the four-object invariant mass, $m_{\mu\mu hh}$, distributions for the data in the CR and a reference signal in the four categories separately, after the event selection (as shown in Table 5.8) is applied. For illustration, an exponential fit on the CR is performed to model the background in the SR. Figure 5.33 shows the $\langle m_{\mu\mu, hh} \rangle$ distribution for the data in the loose CR and two reference signals in the four categories, under the charged kaon assumption for the final state hadrons, after the full selection is applied. The background distri-

Category	leading μ	subleading μ	leading h	subleading h
<i>prompt</i>	≤ 0.3	≤ 0.6	≤ 0.6	≤ 0.8
<i>displaced $\mu\mu$</i>	≤ 0.5	≤ 0.8	≤ 1	≤ 1.6
<i>displaced hh</i>	≤ 0.5	≤ 0.8	≤ 1	≤ 1.6
<i>displaced</i>	≤ 0.5	≤ 0.8	≤ 1	≤ 1.6

Table 5.7: PF Relative isolation thresholds applied in the analysis.

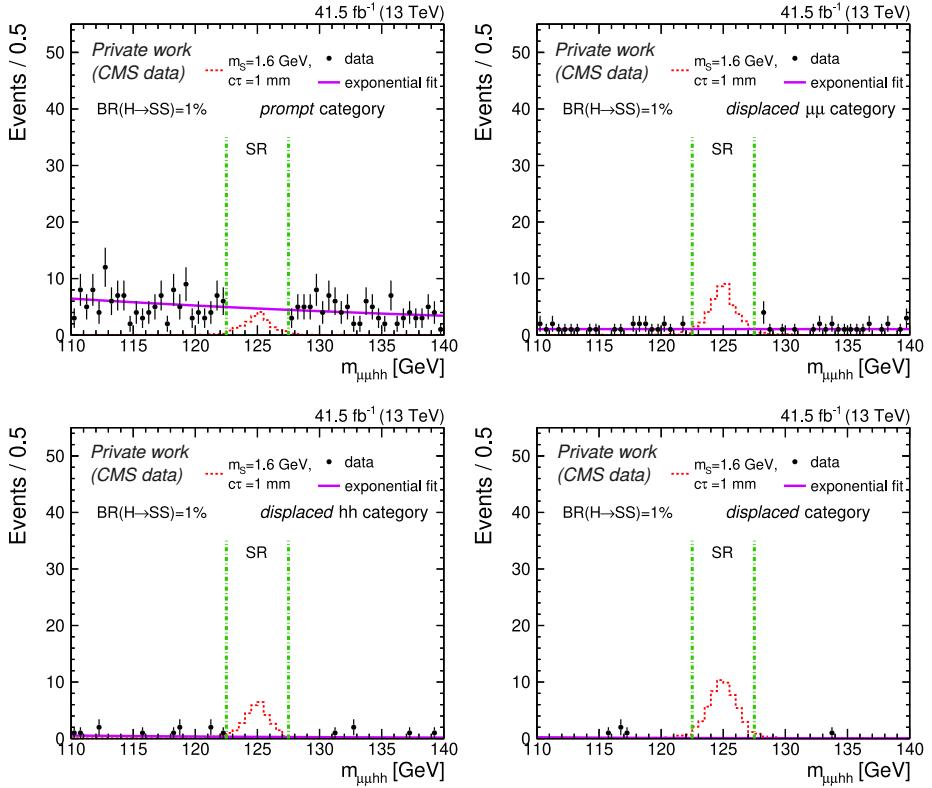


Figure 5.32: Four-object invariant mass distributions for blinded 2017 data (black dots) and for a BSM scalar boson signal of proper lifetime $c\tau = 1$ mm and $m_S = 1.6$ GeV (red histogram) for the *prompt* (upper left), *displaced $\mu\mu$* (upper right), *displaced hh* (lower left) and the *displaced* category (lower right). The event selection criteria, as listed in Table 5.8, have been applied. The green dash-dotted vertical lines delineate the SR.

bution (data in loose CR) shown here is an initial estimate and is limited by low statistics, especially in the non-*prompt* categories. A more robust background estimation strategy is discussed in the next section.

From Fig. 5.33, the *prompt* category already suggests that the expected background events are spread over different values of $\langle m_{\mu\mu, hh} \rangle$ while the signal is concentrated in narrow windows. This motivates the search to be performed in sliding windows of $\langle m_{\mu\mu, hh} \rangle$ around each considered signal mass hypothesis. A Gaussian function is fit to the signal $\langle m_{\mu\mu, hh} \rangle$ distribution to determine the signal mass resolution, σ . The resolution varies with the considered mass hypotheses and is evaluated for each m_S using samples with the proper lifetime $c\tau = 1$ mm. Mass windows of $\pm 2.5\sigma$ are then defined around the mass hypotheses to estimate the signal and background yield in the SR for each event category. The expected signal yield in the SR is taken from simulation with several corrections applied to account for discrepancies between data and simulation. This includes corrections for the pileup reweighting, trigger efficiency, muon reconstruction efficiency, and secondary vertex reconstruction efficiency. These corrections are described in detail in Section 5.7.

Selection	Requirements	Additional Information
Muons	$n_\mu \geq 2, p_T^{\mu_1} \geq 26 \text{ GeV},$ $p_T^{\mu_2} > 5 \text{ GeV}, \eta_\mu < 2.4$	$p_T^{\mu_1} \geq 29 \text{ GeV (2017)},$ $p_T^{\mu_2} > 26 \text{ GeV (2018 di}\mu\text{ trigger)}$
Hadrons	$n_h \geq 2, p_T^h > 5 \text{ GeV}, \eta_h < 2.4,$ $h = \pi^\pm \text{ (h = K}^\pm \text{ for } m_S \geq 1.1 \text{ GeV)}$	
Dimuon	$\Delta R_{\mu\mu} < 0.4, \text{ valid vertex}$	
Dihadron	$\Delta R_{hh} < 0.4, \text{ valid vertex}, p_T^{hh} > 20 \text{ GeV}$	
Loose invariant mass	$m_{\mu\mu} < 5 \text{ GeV}, m_{hh} < 5 \text{ GeV},$ $m_{\mu\mu hh} \in [110, 140] \text{ GeV}$	SR and CR in $m_{\mu\mu hh}$
Di-object invariant mass	$m_{\mu\mu} \sim m_{hh}$	2D diagonal mass cut
Categories	<i>prompt</i> ($L_{xy}^{\mu\mu} / \sigma_{xy}^{\mu\mu} < 40, L_{xy}^{hh} / \sigma_{xy}^{hh} < 30$), <i>displaced</i> $\mu\mu, hh$, <i>displaced</i>	Non- <i>prompt</i> categories from inverting L_{xy} / σ_{xy} alternatively
	<i>prompt</i>	category-wise cuts
Relative isolation	$I_{rel}^{\mu_1} \leq 0.3, I_{rel}^{\mu_2} \leq 0.6, I_{rel}^{h_1} \leq 0.6, I_{rel}^{h_2} \leq 0.8$, $I_{rel}^{\mu_1} \leq 0.5, I_{rel}^{\mu_2} \leq 0.8, I_{rel}^{h_1} \leq 1, I_{rel}^{h_2} \leq 1.6$	$\mu_1, \mu_2 = \text{leading, subleading } \mu$ $h_1, h_2 = \text{leading, subleading } h$

Table 5.8: Summary of the event selection requirements for the analysis. The set of criteria up to (and including) the loose invariant mass row defines the baseline selection. The subscript 1 and 2 refer to the leading and subleading object, respectively.

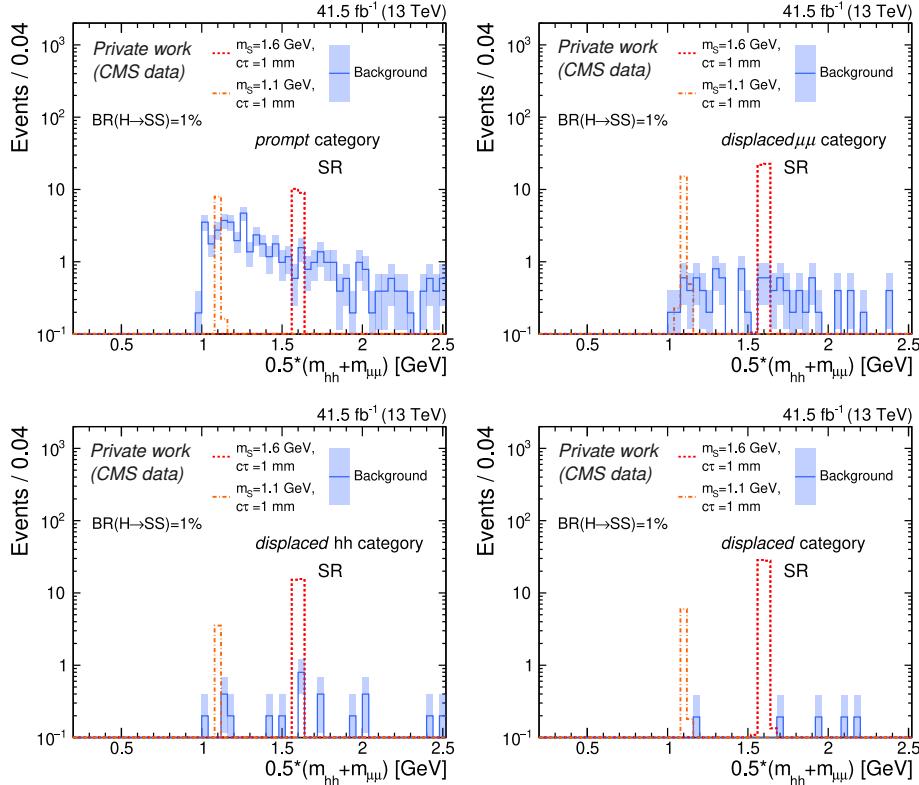


Figure 5.33: Average di-object invariant mass distributions in the SR for the *prompt* (upper left), *displaced* $\mu\mu$ (upper right), *displaced* hh (lower left) and the *displaced* category (lower right). The background prediction, shown as blue histograms, is taken from the sidebands of the $m_{\mu\mu hh}$ distribution of Fig. 5.32. Two signal samples are also presented for BSM scalar bosons of $c\tau = 1$ mm and $m_S = 1.1$ and 1.6 GeV.

5.5 Background estimation

Previously, the expected background in the SR was shown with the data in the sidebands of the $m_{\mu\mu hh}$ distribution (CR and tight CR), while being inclusive in $\langle m_{\mu\mu, hh} \rangle$. Given that the search will be performed in narrow windows around the BSM scalar boson mass hypothesis, the background becomes much smaller. As an example, Fig. 5.34 shows the four-object mass distribution for the 2017 data in the *prompt* category when the $\langle m_{\mu\mu, hh} \rangle$ window for $m_S = 1.6$ GeV is imposed. This then substantially reduces the background as compared to Fig. 5.32 (top left). Similar to before, an exponential fit is performed on the data in the CR to estimate the transfer factor as $f_{\text{SR/CR}} = N_{\text{fit}}^{\text{SR}} / N_{\text{fit}}^{\text{CR}}$, where $N_{\text{fit}}^{\text{SR}}$ ($N_{\text{fit}}^{\text{CR}}$) is the number of events in the SR (CR) extracted from the fit. The events from the CR are then reweighted with the transfer factor to estimate the background contribution in the SR.

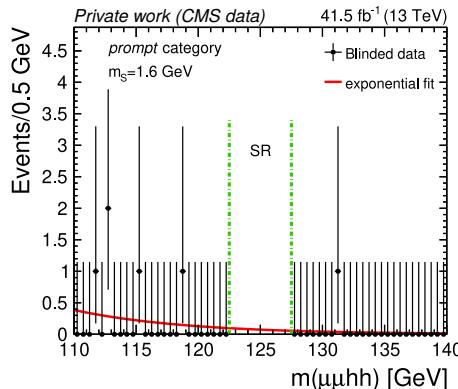


Figure 5.34: The four-object invariant mass distribution for the 2017 data is shown in the *prompt* category after the application of the event selection listed in Table 5.8 and the mass window for $m_S = 1.6$ GeV. The dots represent the data. The red curve is the result of the exponential fit to the CR. The green dash-dotted vertical lines delineate the SR.

The above strategy however, is applicable only to the *prompt* category, since the non-*prompt* categories contain very few events in the CR and provide unreliable fits when the $\langle m_{\mu\mu, hh} \rangle$ window is imposed. This can be inferred from the $\langle m_{\mu\mu, hh} \rangle$ distribution for the CR data in Fig. 5.33 after the full selection is

applied, where the CR data is sparsely distributed. As a result, a relaxed CR is defined for the non-*prompt* categories by removing some of the selection criteria before performing the background estimation in the $\langle m_{\mu\mu, hh} \rangle$ windows.

This proceeds over the following steps. The first step is merging the *displaced* $\mu\mu$, *displaced* hh and the *displaced* categories (shown in Fig. 5.32 top right, bottom left, bottom right) into one common non-prompt category. At this stage, the factor associated to the fraction of events of any particular category that reside in the new merged nonprompt category is calculated as $f_1 = \frac{N_{\text{category}}}{N_{\text{merged category}}}$. Next, the isolation requirements are removed. To estimate the transfer factor between the regions with and without the isolation requirements, the two corresponding $\langle m_{\mu\mu, hh} \rangle$ distributions are compared for the background events lying in the sidebands of the $m_{\mu\mu, hh}$ distribution. Figure 5.35 shows the $\langle m_{\mu\mu, hh} \rangle$ distribution for the merged nonprompt category without the isolation requirements (left) and with the isolation requirements (right) under the charged pion assumption (top) and the charged kaon assumption (bottom) for the final state hadrons. A fit using a Landau function is performed on the $\langle m_{\mu\mu, hh} \rangle$ for the case with no isolation applied². The same function with free normalization is fitted to the event distribution where the isolation criteria are applied, to obtain the transfer factor $f_2 = \frac{N_{\text{isolation, merged}}}{N_{\text{no isolation, merged}}}$. The transfer factor $f_{\text{SR/CR}} = N_{\text{fit}}^{\text{SR}} / N_{\text{fit}}^{\text{CR}}$ for this relaxed CR is then estimated with an exponential fit.

In the method described above, a primary assumption is that the $\langle m_{\mu\mu, hh} \rangle$ distribution remains mostly unchanged for each relaxed cut. Starting with the case of transfer factor f_2 , it is assumed that the $\langle m_{\mu\mu, hh} \rangle$ distribution has the same shape before and after the isolation is applied on the merged non-prompt category. This was tested by comparing a freely floating Landau fit on the $\langle m_{\mu\mu, hh} \rangle$ distribution (isolation criteria applied) with the fit shown in Fig. 5.36. The comparison is shown under the charged kaon assumption (left) and the charged pion assumption (right) for the final state hadrons in Fig. 5.36. The freely floating Landau fit parameters on the isolation applied distribution is found to be compatible with the fit shown in Fig. 5.35 (right).

²The Landau function $f(x; \lambda, \sigma) = \frac{1}{\sigma} L\left(\frac{x-\lambda}{\sigma}\right)$ where $L(z)$ is a standard Landau distribution [152], λ is the most probable value and σ is the width. The standard Landau distribution is defined as $L(z) = \frac{1}{\pi} \int_0^\infty \exp[-t \ln t - zt] \sin(\pi t) dt$.

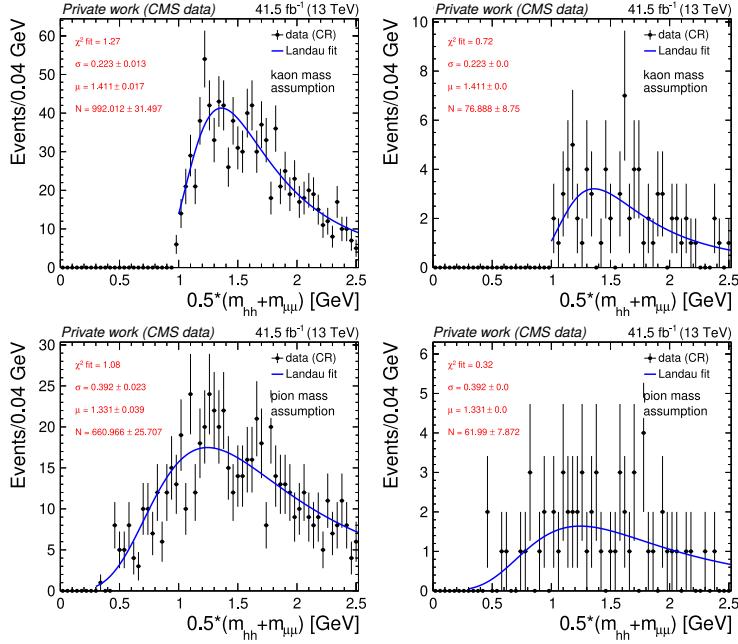


Figure 5.35: Average di-object invariant mass distributions in the merged non-*prompt* category without (left) and with (right) the isolation cuts for the 2017 data. The top row contains the distributions when the charged hadrons in the final state are assumed to be kaons. The bottom row considers the charged pion assumption for the final state hadrons. The data is in the control region. The Landau function with freely floating parameters is fit on the left and the parameters are then used to perform the fit on the right with floating normalization.

For the transfer factor f_1 , it is assumed that it is constant for the various $\langle m_{\mu\mu, hh} \rangle$ windows in which the search is performed. This is a valid assumption if the shape of $\langle m_{\mu\mu, hh} \rangle$ distribution remains the same between the individual non-prompt categories and the merged category. In such a case, f_1 would depend only on the total number of events in the merged and the individual nonprompt categories. Since it is already established that isolation criteria do not change the shape of the $\langle m_{\mu\mu, hh} \rangle$ distribution, Landau fits are performed on the $\langle m_{\mu\mu, hh} \rangle$ distribution for each of the non-prompt categories with freely floating parameters and fixed parameters from the fit on the merged non-*prompt* category (Fig. 5.35 (left)). This is shown in Fig. 5.37 for the charged kaon assumption of the final

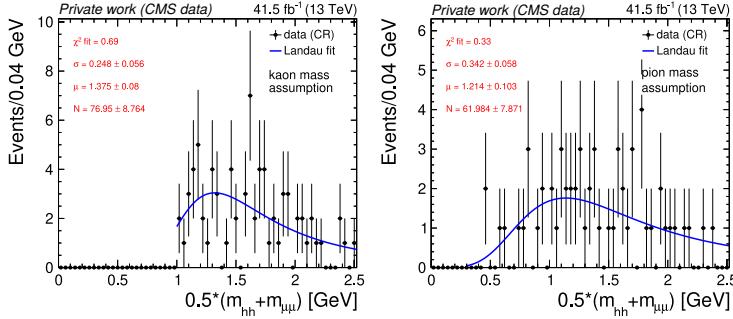


Figure 5.36: Average di-object invariant mass distributions in the merged non-*prompt* category with the isolation cuts for the 2017 data under the charged kaon assumption (left) and the charged pion assumption (right) for the final state hadrons. The data is in the control region. The Landau function with freely floating parameters is used to model the distribution.

state hadrons. It is seen that the fits between the fixed and freely floating parameters are mostly compatible within the uncertainties on the parameters. A similar observation is also present for the results with the charged pion assumption for the final state hadrons.

To summarize, the background yield for the non-*prompt* categories is measured by weighting the CR events with f_1, f_2 and measuring the transfer factor $f_{\text{SR/CR}}$ after applying the $\langle m_{\mu\mu, hh} \rangle$ window according to the signal mass hypothesis. For the *prompt* category, only the transfer factor $f_{\text{SR/CR}}$ is applied. While this procedure is performed for each era of the Run 2, the combined results for the full Run 2 is shown in the next subsection to provide an overview of the background estimation, signal and the observed data.

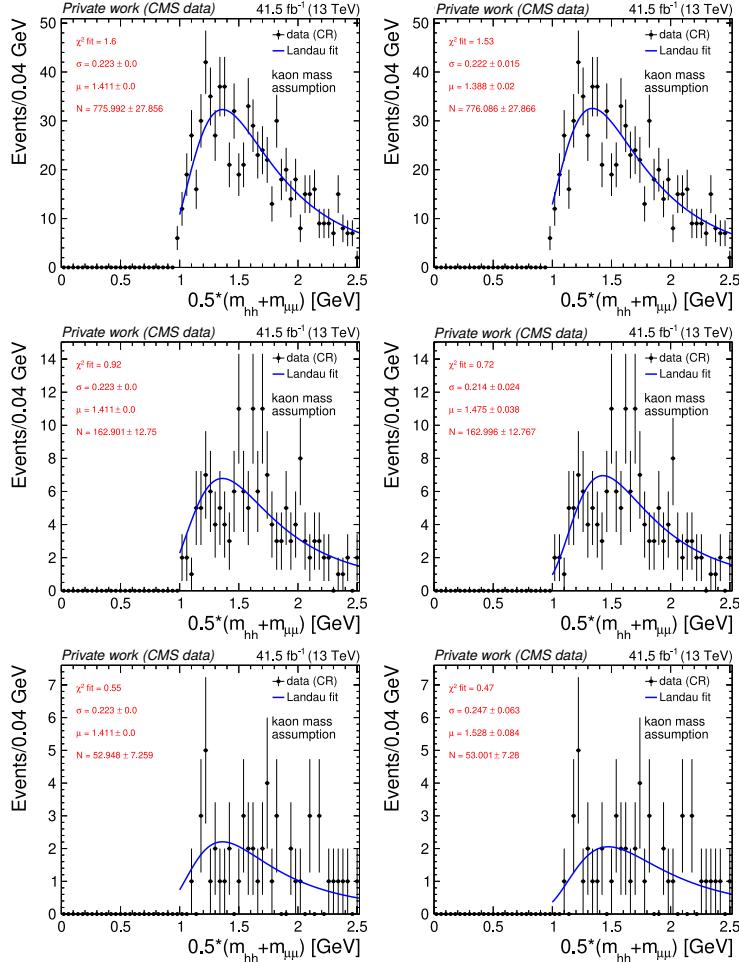


Figure 5.37: Average di-object invariant mass distributions in the non-*prompt* categories, *displaced* $\mu\mu$ (top), *displaced* hh (middle), *displaced* (bottom), without the isolation cuts for the 2017 data under the charged kaon assumption for the final state hadrons. The data is in the control region. A Landau fit is performed with fixed parameters from the fit in Fig. 5.35 and floating normalization on the left column while the right column has freely floating parameters.

5.6 Full Run 2 data and unblinding

The previous sections described the analysis development and the strategy while avoiding any bias from data by blinding the signal region. In this section, a first look at the unblinded results is shown for the full Run 2 dataset to provide a qualitative look at the results of the analysis. The two following sections present the systematic corrections and uncertainties (Section 5.7) and the results of the analysis with the expected and observed limits(Section 5.8).

Figure 5.38 and 5.39 show the $m_{\mu\mu\text{hh}}$ distribution for the full Run 2 dataset after unblinding the signal region and application of all the event selection cuts (listed in Table 5.8) under the charged kaon and charged pion assumption for the final-state hadrons, respectively. Two signal samples, $m_S = 0.6$ GeV and $m_S = 1.6$ GeV with $c\tau = 1$ mm are chosen as references. While the background distribution is similar in both cases (expected), the signal yield for $m_S = 1.6$ GeV is higher as compared to $m_S = 0.6$ GeV since the lighter BSM scalar boson masses produce more collimated and displaced decays, resulting in a comparatively lower reconstruction efficiency.

Figure 5.40 shows the corresponding distribution for the merged nonisolated non-*prompt* category. The relaxed CR events, weighted with f_1, f_2 and $f_{\text{SR/CR}}$ produces the background prediction for the $\langle m_{\mu\mu,\text{hh}} \rangle$ distribution shown as the blue histograms in Fig. 5.41 and 5.42 for each of the non-*prompt* categories. The background prediction in the *prompt* category is taken from the CR events (Fig. 5.38 and 5.39 top left) after passing all selection criteria and weighted with $f_{\text{SR/CR}}$. It is observed that the data is in agreement with the background estimation within statistical fluctuations.

Finally, the expected background and signal yields are extracted after applying the $\langle m_{\mu\mu,\text{hh}} \rangle$ window corresponding to each signal mass hypothesis. As an example, Fig. 5.43 shows the $m_{\mu\mu\text{hh}}$ distribution in the *prompt* category (left) and the merged non-*prompt* categories (right), for $\langle m_{\mu\mu,\text{hh}} \rangle = 0.6$ GeV (top) and for $\langle m_{\mu\mu,\text{hh}} \rangle = 1.6$ GeV (bottom). The merged non-*prompt* category shows the unweighted distribution without the transfer factors for events passing the relaxed criteria discussed before. Table 5.9 shows the signal yield for $m_S = 0.6$ GeV and $c\tau = 1$ mm (top), and $m_S = 1.6$ GeV and $c\tau = 1$ mm (bottom), the

background prediction and the observed data. In some cases, for instance for $m_S = 1.6$ GeV, the background yield is under-predicted as compared to the observation and seems to lie beyond the $\pm 1\sigma$ statistical uncertainties quoted in the Table. However, additional systematic uncertainties need to be taken into account when determining the final limits. These uncertainties are discussed in the next section.

Region	$m_S = 0.6$ GeV, $c\tau = 1$ mm	Background prediction	Data
Prompt	48 ± 2 (stat.) ± 7 (syst.)	2.0 ± 0.5	2
Displaced $\mu\mu$	57 ± 1 (stat.) ± 14 (syst.)	0.15 ± 0.04	0
Displaced hh	39 ± 1 (stat.) ± 7 (syst.)	0.08 ± 0.02	0
Displaced	45 ± 1 (stat.) ± 11 (syst.)	0.010 ± 0.003	0

Region	$m_S = 1.6$ GeV, $c\tau = 1$ mm	Background prediction	Data
Prompt	111 ± 3 (stat.) ± 17 (syst.)	4.7 ± 0.9	1
Displaced $\mu\mu$	187 ± 4 (stat.) ± 46 (syst.)	1.4 ± 0.1	0
Displaced hh	121 ± 3 (stat.) ± 22 (syst.)	0.52 ± 0.05	0
Displaced	172 ± 3 (stat.) ± 42 (syst.)	0.15 ± 0.02	0

Table 5.9: Number of observed events and background predictions for the full Run 2 dataset in the SR under the BSM scalar boson mass hypothesis, $m_S = 0.6$ GeV (top) and $m_S = 1.6$ GeV (bottom). In comparison, the signal yield expected in the full Run 2 dataset, assuming $\mathcal{B}(H \rightarrow SS) = 1\%$, is quoted.

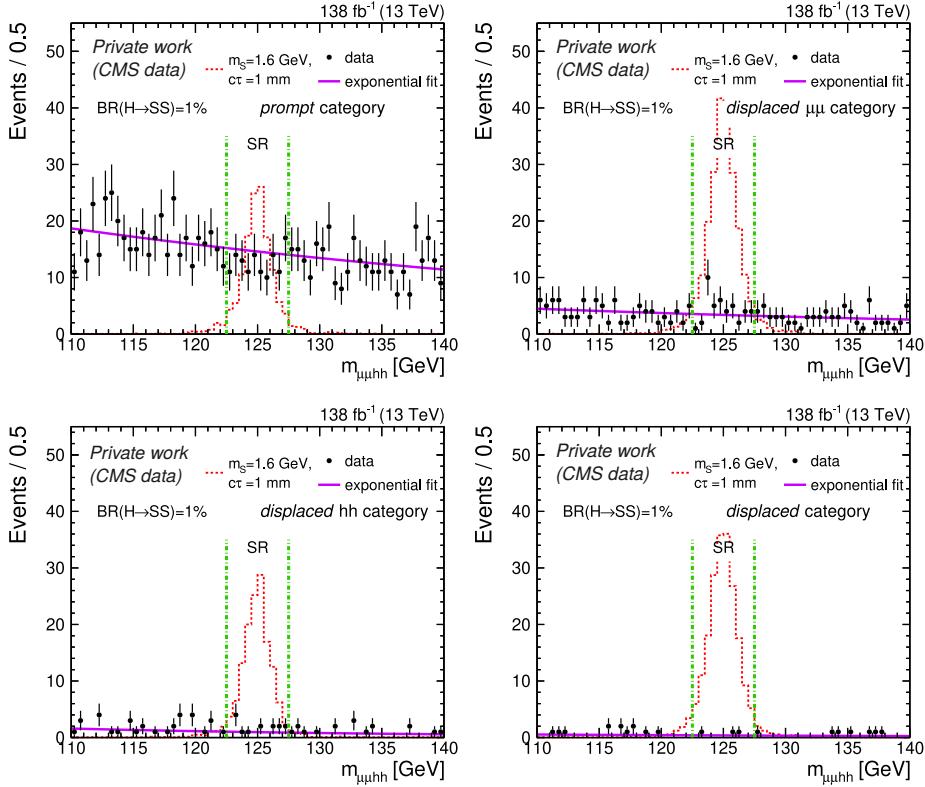


Figure 5.38: Four-object invariant mass distributions for full Run 2 data (black dots) and for a BSM scalar signal of proper lifetime $c\tau = 1$ mm and $m_S = 1.6$ GeV (red histogram) for the *prompt* (upper left), *displaced* $\mu\mu$ (upper right), *displaced* hh (lower left) and the *displaced* category (lower right). The event selection criteria, as listed in Table 5.8, have been applied. The green dash-dotted vertical lines delineate the SR.

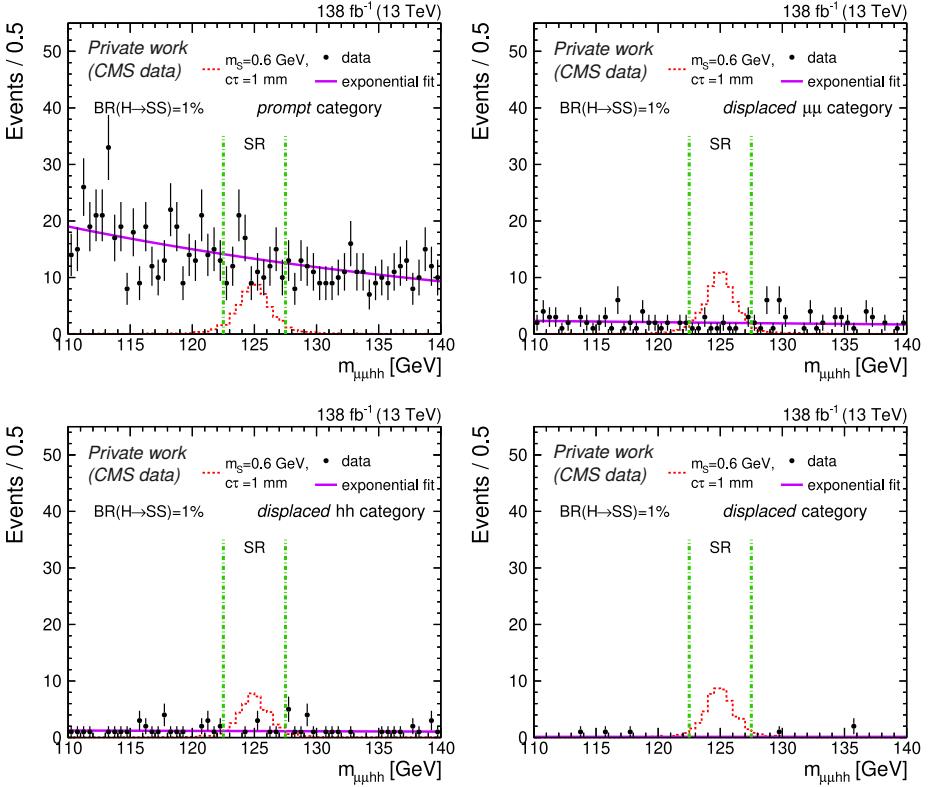


Figure 5.39: Four-object invariant mass distributions for full Run 2 data (black dots) and for a BSM scalar signal of proper lifetime $c\tau = 1$ mm and $m_S = 0.6$ GeV (red histogram) for the *prompt* (upper left), *displaced* $\mu\mu$ (upper right), *displaced* hh (lower left) and the *displaced* category (lower right). The event selection criteria, as listed in Table 5.8, have been applied. The green dash-dotted vertical lines delineate the SR.

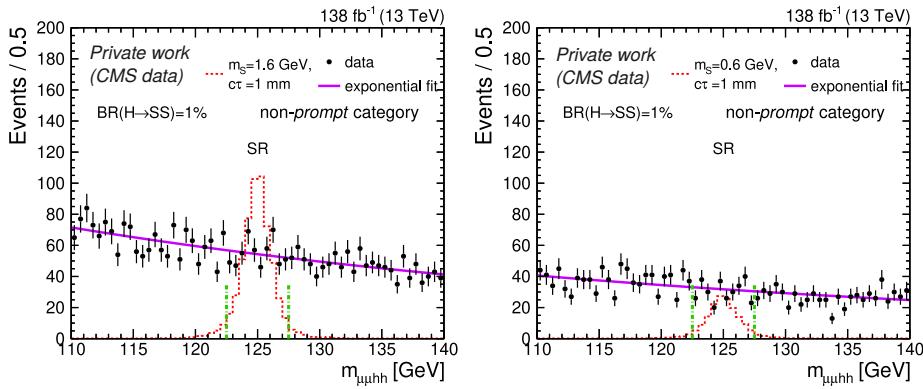


Figure 5.40: Four-object invariant mass distributions for full Run 2 data (black dots) in the merged nonisolated non-*prompt* category (relaxed CR) under the charged kaon assumption (left) and the charged pion assumption (right) for the final-state hadrons. The BSM scalar boson signals with $m_S = 0.6$ GeV, $c\tau = 1$ mm (red histogram, left) and $m_S = 1.6$ GeV, $c\tau = 1$ mm (red histogram, right) are shown for reference with $\mathcal{B}(H \rightarrow SS) = 1\%$. The green dash-dotted vertical lines delineate the SR. The distributions are not weighted by the relevant transfer factors.

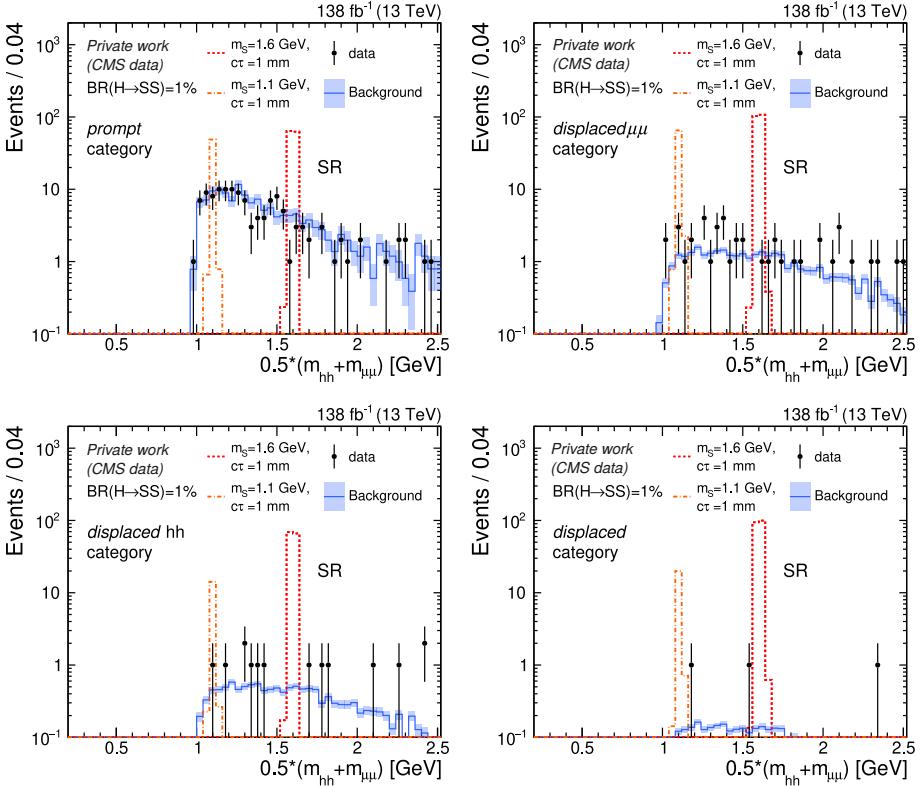


Figure 5.41: Average diobject invariant mass distributions in the SR for the *prompt* (upper left), *displaced* $\mu\mu$ (upper right), *displaced* hh (lower left) and the *displaced* category (lower right) under the charged kaon assumption for the final-state hadrons. The event selection criteria of Table 5.8 have been applied. The observed number of events in the full Run 2 dataset are shown as black dots and the background prediction (estimated from the events in the CR as explained in Section 5.5) as blue histograms. Two signal samples are also presented for BSM scalar bosons of $c\tau = 1$ mm and $m_S = 1.1$ and 1.6 GeV as reference with $\mathcal{B}(H \rightarrow SS) = 1\%$.

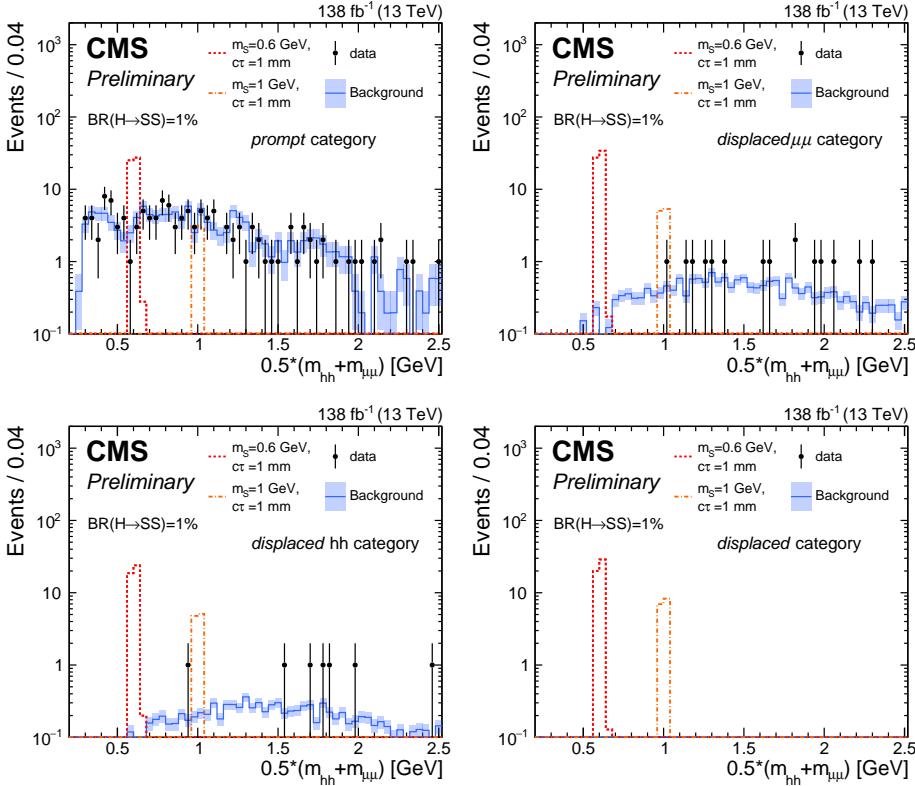


Figure 5.42: Average diobject invariant mass distributions in the SR for the *prompt* (upper left), *displaced* $\mu\mu$ (upper right), *displaced* hh (lower left) and the *displaced* category (lower right) under the charged pion assumption for the final-state hadrons. The event selection criteria of Table 5.8 have been applied. The observed number of events in the full Run 2 dataset are shown as black dots and the background prediction (estimated from the events in the CR as explained in Section 5.5) as blue histograms. Two signal samples are also presented for BSM scalar bosons of $c\tau = 1$ mm and $m_S = 0.6$ and 1 GeV as reference with $\mathcal{B}(H \rightarrow SS) = 1\%$. Figures shown in Ref. [153].

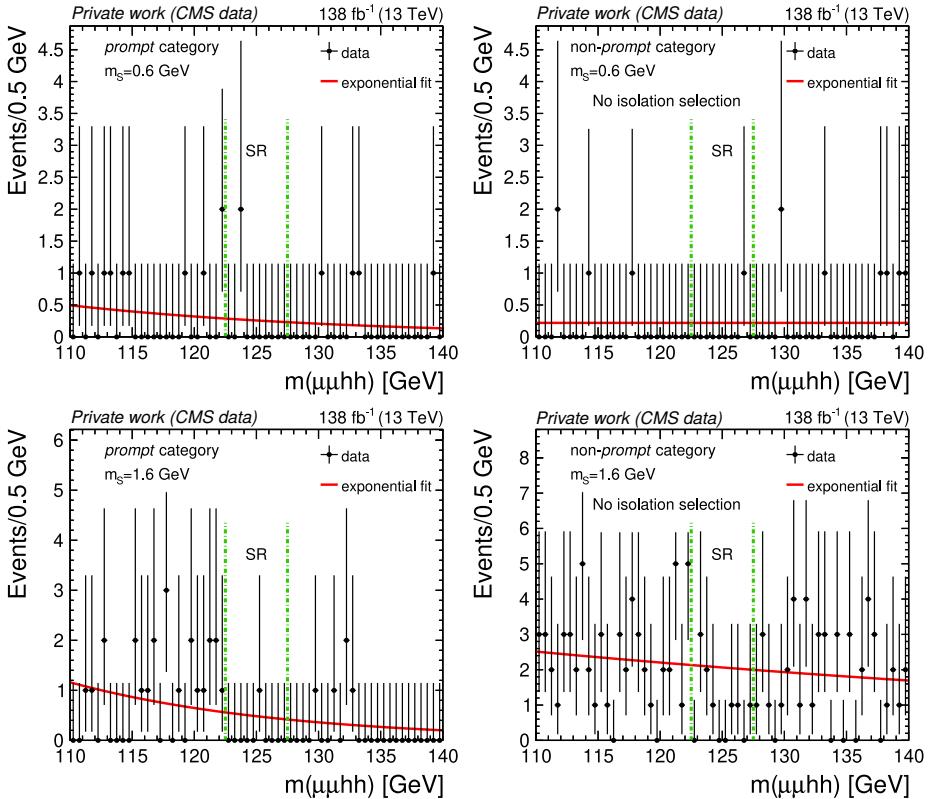


Figure 5.43: The four-object invariant mass distribution for the full Run 2 data set is shown in the *prompt* (left) and merged nonisolated *non-prompt* (right) categories after the application of the event selection listed in Table 5.8 and the mass window for $m_S = 0.6$ GeV (top) and $m_S = 1.6$ GeV. The dots represent the data. The red curve is the result of the exponential fit to the CR. The green dash-dotted vertical lines delineate the SR. The distributions for the merged nonisolated *non-prompt* category are not weighted by the relevant transfer factors.

5.7 Systematic Corrections and Uncertainties

This section presents the systematic uncertainties and the relevant corrections that have been applied in the analysis. In particular, a few dedicated corrections were carried out to describe the specific signal kinematics. These are described in more detail as compared to the standard corrections that are applied across various CMS analyses.

5.7.1 Background Estimation Uncertainties

Statistical Uncertainty

The statistical uncertainty on the background estimate depends on the number of events in the control region. For the *prompt* category, the total number of events in the control region ranges from 0–20 events in each data-taking era, across the different BSM scalar boson mass hypotheses. For the non-*prompt* categories, this ranges from 0–50 events in each era. In general, uncertainties during the limit setting procedure are modelled using the log-normal distribution since they typically come from calibration measurements [154]³. However, due to limited statistics in the background, the associated uncertainty is modelled using a gamma distribution which is well-suited to describe the error due to a limited sample size⁴. In this case, the relative error is equal to $1/\sqrt{N+1}$ where N is the number of events in the CR (relaxed CR) for *prompt* category (non-*prompt* categories). Table 5.10 and 5.10 show this statistical error for the different categories, BSM scalar boson mass hypotheses and the data-taking eras. The statistical uncertainty shown here has the dominant impact in the overall uncertainty in the measurement of the final sensitivity of the search.

Systematic Uncertainty

Across all categories, it is assumed that the four-object invariant mass ($m_{\mu\mu\text{hh}}$) distribution of the background events follows an exponential nature. To study

³ $X \sim \text{lognormal}(\mu, \sigma)$ if $\ln X \sim \mathcal{N}(\mu, \sigma^2)$.

⁴Gamma function is defined as $f(x; \alpha, \lambda) = \frac{x^{\alpha-1} \exp^{-\lambda x} \lambda^\alpha}{\Gamma(\alpha)}$ where α is the shape parameter and λ is the rate parameter.

m_S / year	early 2016	late 2016	2017	2018
0.4 GeV	71%	41%	29%	41%
0.6 GeV	58%	50%	50%	32%
0.8 GeV	58%	71%	28%	28%
0.9 GeV	50%	58%	24%	45%
1 GeV	71%	58%	35%	27%
1.1 GeV	58%	35%	33%	24%
1.2 GeV	50%	38%	29%	24%
1.4 GeV	50%	58%	41%	30%
1.6 GeV	71%	50%	38%	23%
1.8 GeV	100%	71%	41%	41%
2 GeV	100%	100%	41%	45%

Table 5.10: Relative statistical uncertainty error on the background estimation with respect to the BSM scalar boson masses for *prompt* category.

m_S / year	early 2016	late 2016	2017	2018
0.4 GeV	58%	100%	100%	58%
0.6 GeV	100%	71%	38%	31%
0.8 GeV	58%	38%	32%	25%
0.9 GeV	45%	58%	26%	22%
1 GeV	58%	71%	19%	24%
1.1 GeV	33%	35%	21%	19%
1.2 GeV	29%	41%	18%	17%
1.4 GeV	30%	27%	18%	14%
1.6 GeV	33%	58%	16%	14%
1.8 GeV	35%	58%	16%	16%
2 GeV	45%	32%	22%	16%

Table 5.11: Relative statistical uncertainty error on the background estimation with respect to the BSM scalar boson masses for non-*prompt* categories.

the systematic effect of this choice, alternative functional forms, the power law and linear functions, to the exponential model were considered. The variation in the resulting background predictions, propagated via the transfer factor $f_{\text{SR/CR}}$, were found to be very small as compared to the statistical uncertainty and were neglected.

For the non-*prompt* categories, additional sources of uncertainties are in-

cluded to account for the multiple transfer factors. The uncertainty on these transfer factors are described below:

- f_1 uncertainty: This is estimated by propagating $\sigma = \sqrt{N}$ on $N(\text{category})$ and $N(\text{merged category})$ for each of the 3 non-*prompt* categories. Tables 5.12, 5.13 show the uncertainties for each era for the cases where $h^\pm = \pi^\pm$ and $h^\pm = K^\pm$ in the final state, respectively. This results in uncertainties of up to 100% for the *displaced* category which arises from nearly zero events in the sidebands of $m_{\mu\mu hh}$ distribution after applying the full event selection.
- $f_2 = N_{\text{iso}}^{\text{scalar mass}} / N_{\text{no iso}}^{\text{scalar mass}}$ uncertainty: This is estimated by propagating the errors for $N_{\text{isolation,merged}}$ and $N_{\text{no isolation,merged}}$ as obtained from the fit on the $\langle m_{\mu\mu, hh} \rangle$ distribution. Table 5.14 shows the uncertainties for each era for the cases where $h^\pm = \pi^\pm$ and $h^\pm = K^\pm$ in the final state. The uncertainty ranges from 10–35%.

category	early 2016	late 2016	2017	2018
<i>displaced</i> $\mu\mu$	30%	30%	16%	13%
<i>displaced</i> hh	61%	45%	22%	17%
<i>displaced</i>	100%	100%	39%	29%

Table 5.12: Uncertainty on f_1 on background for different non-*prompt* categories with $h^\pm = K^\pm$.

category	early 2016	late 2016	2017	2018
<i>displaced</i> $\mu\mu$	47%	38%	18%	16%
<i>displaced</i> hh	64%	44%	22%	19%
<i>displaced</i>	76%	100%	51%	59%

Table 5.13: Uncertainty on f_1 on background for different non-*prompt* categories with $h^\pm = \pi^\pm$.

	early 2016	late 2016	2017	2018
$h^\pm = K^\pm$	22%	22%	12%	9%
$h^\pm = \pi^\pm$	33%	27%	13%	11%

Table 5.14: Uncertainty on f_2 on background for different non-*prompt* categories with $h^\pm = K^\pm$ and $h^\pm = \pi^\pm$.

5.7.2 Signal Uncertainties

Luminosity

The uncertainty in the integrated luminosity measurement is taken according to the recommendations from the [142, 155, 156] for each era. This varies from 1–3% with the largest uncertainty for the 2018 data-taking.

Pileup reweighting

The pileup profile in simulation is reweighted to match the profile observed in data by applying weights based on the number of vertices in the event. The pileup reweighting uncertainty is measured from $\pm 1\sigma$ variation in pp inelastic cross section and measuring the change in the signal yield for the different BSM scalar boson masses considered. The size of the pileup uncertainty was estimated to be about $\sim 1\%$ for the signal masses and lifetimes.

Reconstruction of close muons

One of the complementary CMS analyses that targets similar light BSM particles as this analysis is Ref. [64], as discussed in Section 1.2.3. The analysis in Ref. [64] performs a model-independent search for new bosons within the mass range $0.21 < m_a < 60$ GeV, targeting muons produced in the boson decays. Since the light bosons lead to a highly boosted and collimated pair of muons, a dedicated study measuring the effect related to tracking for low-mass dimuons at high momentum ($m_a \in [0.25, 5]$ GeV) was performed. The systematic uncertainty for this effect was estimated to be around 1.2% for each dimuon in the event. Since this analysis targets similar physics and kinematic phase space, the same uncertainty is assigned here to account for the reconstruction effect i.e. 1.2% for the dimuon in the final state.

Muon ID

The corrections for the performance of the loose muon identification in simulation and data are taken from the central measurements by the CMS experiment [126]. A conservative estimate of the uncertainty on this is taken as 0.6% per muon (1.2% in total), following the approach of Ref. [64].

Reconstruction of displaced muons

A search for long-lived particles decaying to a pair of muons was performed in Ref. [122] where a dedicated study with cosmic muons was performed to account for the uncertainty on the effect of displacement on muon reconstruction. A systematic uncertainty of 5% per muon was assigned to account for mismodelling of displaced muon reconstruction in simulation. In this analysis, we apply the results of that study and assign an uncertainty of 5% per muon (10% in total).

Reconstruction of hadronic tracks

Charged hadrons, classified as either charged pions or charged kaons are used in the final state. To assess the systematic uncertainty associated with the reconstruction of charged hadron tracks, the analysis refers to the study in Ref. [117] which measured the tracking efficiency using D-meson decays. Using the exclusive decay channels, $D^0 \rightarrow K\pi$ and $D^0 \rightarrow K\pi\pi\pi$, in events with $D^* \rightarrow D^0\pi$, the relative efficiency between the two decay modes is measured in data and compared to the values by Particle Data Group [157]. The uncertainty on this measurement is taken as the systematic uncertainty associated to the charged hadrons in the analysis and is between 2–3% per track across the different eras of Run 2.

Muon and hadron isolation

While cuts on the relative isolation were applied on the leading object and the subleading object separately, the hadronic activity surrounding both the objects is similar because of the collimation. Therefore, a correlated cut using the leading and subleading object isolation was chosen, as discussed in Section 5.4.7.

Following the same reasoning, the uncertainty on the cut is applied as a whole for the muon relative isolation cuts and the hadron relative isolation cuts in simulation.

A conservative estimate is taken for the systematic uncertainty to cover any differences between the performance in isolation between simulation and data. The isolation variable is expected to perform worse in data since the hadronic activity in simulation is less compared to data. The uncertainty is taken to be $(1 - \epsilon)/2$ where $\epsilon = N_{\text{pass iso.}}/N_{\text{before iso.}}$ is the signal cut efficiency. In the *prompt* category, the uncertainty associated with the muon (hadron) isolation is 5–10 (5–15)%, as a function of the scalar mass. In the non-*prompt* categories, the isolation requirements are much looser, leading to highly efficient isolation cuts, and the uncertainty in the muon (hadron) isolation leads to a 2.5 (10)% uncertainty in the signal yield.

Secondary vertex reconstruction

In Section 3.2.2, the displaced vertex reconstruction efficiency measurements in Ref. [119] were described briefly. In this analysis, Ref. [119] was extended to derive the systematic corrections and uncertainties associated with secondary vertex reconstruction and the dimuon / dihadron mass resolution. These studies were performed by a fellow colleague from Vrije Universiteit Brussel, and are briefly discussed here.

In this study, $K_S^0 \rightarrow \pi^+ + \pi^-$ decays in Drell-Yan events are selected. Since K_S^0 travels a certain distance before decaying, the decay results in a displaced vertex with two charged pion tracks in the final state. By measuring the number of such reconstructed secondary vertices in simulation and data, a correction and uncertainty is derived as a function of the displacement (L_{xy} or Δ_{2D}) and the reconstructed dihadron p_T .

There are two important extensions that were considered when extending the results of the existing study. The first extension is measuring the corrections with respect to the L_{xy} or Δ_{2D} significance instead of Δ_{2D} since in this analysis, it is the significance that is used as the measure of the displacement of the vertex. The second is the extension of the p_T -range covered in Ref. [119]. Since most of the neutral kaons are produced at low- p_T , the results of the original study were

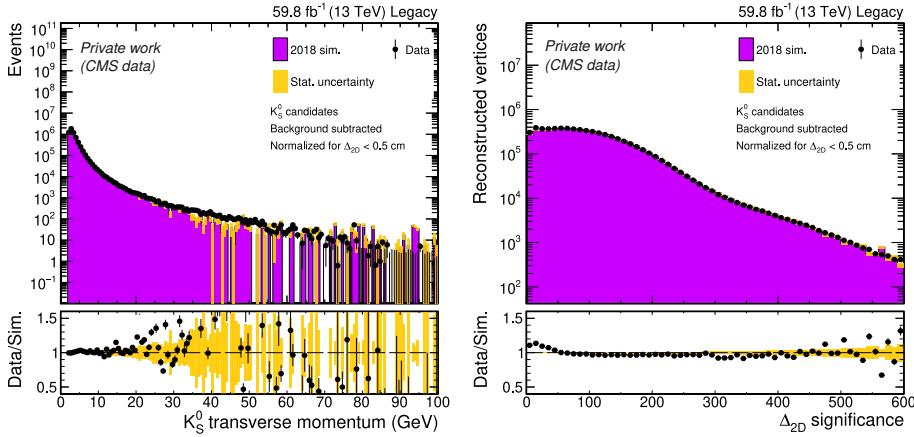


Figure 5.44: Data and simulation plots for the K_S^0 transverse momentum (left) and the L_{xy} significance (right) for the 2018 data-taking era of Run 2.

performed up to 20 GeV in the reconstructed dipion p_T . However, since the event selection in our analysis ensures that the dihadron p_T is at least 20 GeV, the extension of a p_T bin beyond 20 GeV is a vital addition.

The event selection for this study proceeds with the selection $Z \rightarrow \mu\mu$ events with only two well-reconstructed muons in the event. Along with this, requiring the reconstructed dimuon mass to be in a small window around the Z boson mass ensures a selection of well-reconstructed Z boson events. The next step consists of selecting well-reconstructed tracks and forming all possible pairs of oppositely charged pions that share a secondary vertex. The Kalman vertex fitter is used to reconstruct the secondary vertices which are then passed through some quality cuts. Finally, the invariant mass of the track pairs with reconstructed secondary vertices are required to fall in a fixed window centered at the K_S mass ($= 0.493$ GeV). The resulting distribution of neutral kaons selected with this selection is shown in Fig. 5.44. Note that since it is known that the vertex reconstruction for promptly produced particles is $\sim 100\%$, the simulation is normalized to data in the first bin of the Δ_{2D} distribution. The normalization factor is propagated to all events.

The reconstructed K_S mass is binned in p_T and Δ_{2D} -significance in simulation and data. The distribution is then fitted to a sum of Gaussians that provide the

integral or the total number of events in the bin. The ratio of the integrals in data and simulation provide the data-to-simulation scale factors which provide the corrections and the uncertainties associated with it.

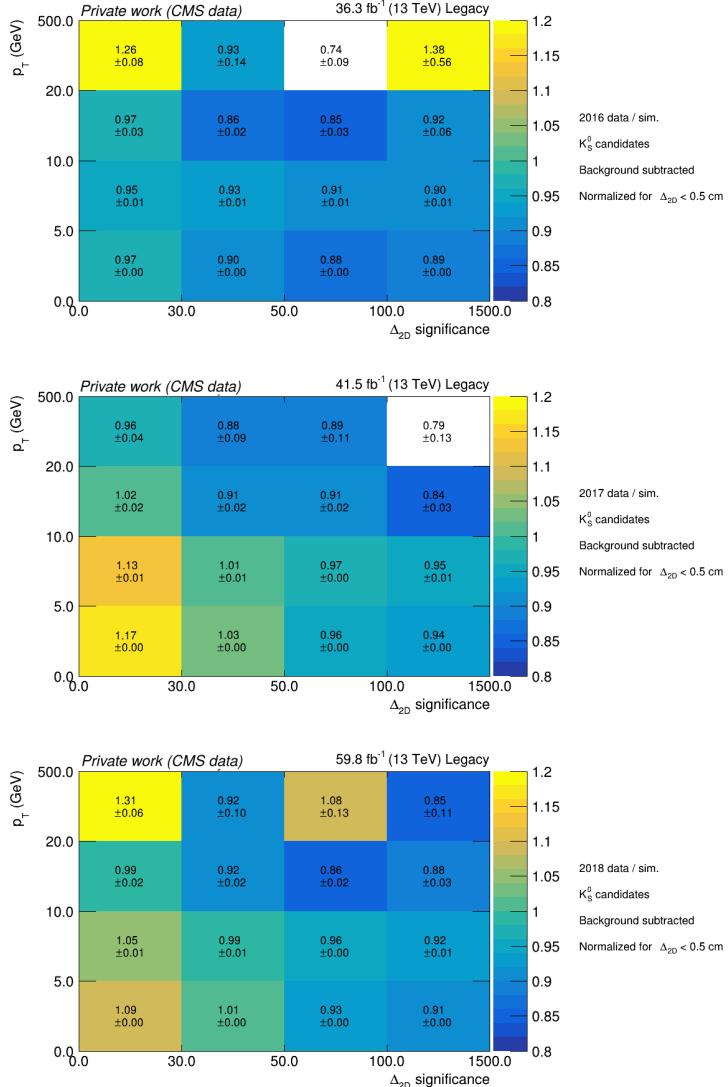


Figure 5.45: Data-to-simulation scale-factors determined for displaced tracking and vertexing of K_S^0 mesons as a function of transverse momentum ($p_T^{\pi\pi}$) and displacement significance ($L_{xy}^{\pi\pi} / \sigma_{xy}$), split per data-taking year.

Fig. 5.45 shows the data-to-simulation scale-factors for the secondary vertex reconstruction performance for different eras of Run 2. The scale-factors derived in this study are then applied to the analysis. The scale-factors are found to be typically within 10% for most bins. As most K_S^0 are produced with low momentum, the study gets statistically limited for $p_T > 20$ GeV and leads to a varied range of uncertainties (5-40% relative uncertainty). Additionally, the low statistics in this bin lead to larger uncertainties when the 2016 dataset is split into the early and late eras. Therefore, the study is carried out for the entire 2016 dataset and the same uncertainties and scale-factors are accordingly applied to both the eras of 2016.

To assign the relative uncertainties, a conservative estimate is taken from the uncertainties present in Fig. 5.45. This is chosen according to the event category and the L_{xy} significance distribution expected in that category. For the *prompt* category, since all events fall within $L_{xy}^{\mu\mu}/\sigma_{xy}^{\mu\mu} < 40, L_{xy}^{hh}/\sigma_{xy}^{hh} < 30$, the uncertainty corresponding to $p_T > 20$ GeV and $L_{xy}/\sigma_{xy} < 30$ is chosen for both the dimuon and dihadron vertex. For the *displaced $\mu\mu$* category, since $L_{xy}^{hh}/\sigma_{xy}^{hh} < 30$, the same uncertainty as for the *prompt* category is chosen for the dihadron vertex. For the dimuon vertex, a conservative uncertainty is chosen by taking the maximum relative uncertainty possible from the bins with $L_{xy}/\sigma_{xy} > 30$. For the other non-*prompt* categories, a similar logic is followed in choosing the uncertainty estimate.

Dimuon/Dihadron mass resolution

The 2-dimensional mass selection, which ensures the scalar mass consistency between the dimuon and the dihadron, depends on their respective mass resolutions from simulation. By extending the study presented in the previous subsection, possible mismodelling of the resolution between simulation & data is accounted for through the related systematic uncertainty.

Following the selection described earlier, the reconstructed K_S mass is binned in p_T and Δ_{2D} -signficance, similar to before. Instead of performing a fit with a sum of Gaussians, a single Gaussian fit is performed on the core of the distribution to extract the resolution of the reconstructed K_S mass in data and simulation.

The extracted resolution in simulation is then used to define the mass bounds

as $\mu + 2.5\sigma^{MC}$. The systematic effect that is targeted here is the number of events in data that are accepted within these mass bounds when the bound is set with respect to σ^{data} instead. Therefore, the ratio $N_{data}^{\pm 2.5\sigma_{data}}/N_{data}^{\pm 2.5\sigma_{MC}}$ provides the offset and the associated uncertainty due to the bound that is derived from simulation.

Figure 5.46 shows this ratio for the different eras of Run 2. It is found to be relatively stable with a correction of $< 5\%$ in all bins with uncertainties of $\sim 1 - 10\%$ except for one bin in 2016 which is statistically limited. Since this correction is small, only the uncertainty is included in the analysis. Similar to the approach before, the assigned uncertainties are chosen according to the category and uses the most conservative estimates as before. Since the mass resolution between the reconstructed dimuon and dihadron are similar (Fig. 5.20), the uncertainties obtained here are accounted for both the dimuon and dihadron mass resolutions.

Trigger

The isolated single muon trigger is a commonly used trigger in the physics analyses involving muons in the final state. As a result, the corrections and uncertainties associated with the trigger efficiency are centrally calculated by the CMS experiment. A data-driven approach called the Tag-and-Probe is often employed for these measurements. The philosophy behind the method is to select known dimuon resonances in data, for instance Z , J/ψ or Υ , and then measure the efficiency of a particular set of conditions on the daughter muons. In particular, the ‘tag’ muon is chosen to pass strict criteria to ensure the selection of good quality events. The ‘probe’ muon is then required to pass a loose set of criteria to capture the pair of muons coming from the resonance. The efficiency is then calculated by measuring the number of probes that pass or fail the quantity being tested, for instance the trigger efficiency, as $\epsilon = \frac{N_{\text{pass}}}{N_{\text{pass}} + N_{\text{fail}}}$. The ratio of the efficiency in data and simulation leads to the correction factor that is applied to simulation. The associated uncertainty in the correction factor is taken as the systematic uncertainty.

The central measurements of the isolated single muon trigger performance are derived using the tag-and-probe method in dimuon decays of the Z boson.

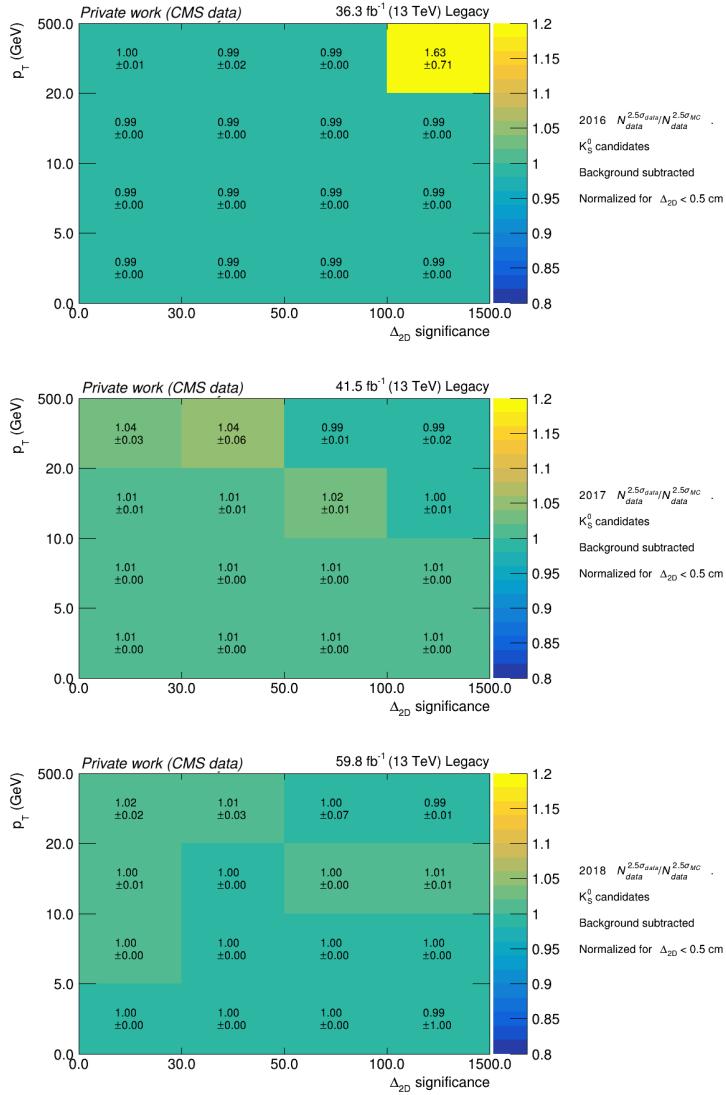


Figure 5.46: Data-to-simulation scale-factors resulting from the (yield @ $\pm 2.5\sigma_{MC}$) vs (yield @ $\pm 2.5\sigma_{data}$) study, given as a function of radial displacement and transverse momentum, split per data-taking year

Taking into account that the signal muon final states in this analysis are much more collimated ($\Delta R < 0.4$ as seen in Fig. 5.5) than what is present in Z decays (dominantly back-to-back decays i.e. $\Delta R \sim 3$), the standard results do not ac-

curately capture the effect of this collimated environment on the isolated single muon trigger efficiency. Moreover, the trigger uses isolation as a selection criteria in its path. As discussed before, the isolation is calculated within a cone of ΔR , which for this analysis, can be affected due to the closeness of the muons.

Accordingly, a dedicated efficiency measurement was undertaken for the isolated single muon trigger where a Tag-and-Probe method utilizing J/ψ was used. J/ψ s provide a similar kinematic domain as they are in a similar mass range as the targeted BSM scalar bosons and decays as a collimated pair of muons. A selection inspired from the muon soft ID studies carried out in Ref. [126] is used and summarized in Table 5.15.

To ensure the selection of an unbiased sample of J/ψ , the prescaled single muon triggers with low p_T thresholds have been chosen. The list of triggers used, shown in Table 5.15, have similar trigger paths with different p_T thresholds and corresponding prescales applied to ensure a low trigger rate. After taking into account the prescale values, the measurement uses datasets with effective luminosities of 0.35 fb^{-1} , 0.06 fb^{-1} , 0.65 fb^{-1} and 0.11 fb^{-1} for early 2016, late 2016, 2017 and 2018 data-taking periods respectively. Furthermore, a requirement to match the tag muon with the triggering object is imposed to ensure that no selection on the probe muons is introduced at this stage. The baseline selection also requires the tag muon to pass the tight ID, $p_T > 8 \text{ GeV}$ and $|\eta| < 2.4$ to ensure a well-reconstructed tag muon.

In addition to this, two new cuts are introduced based on the tag muon tracker isolation ($relIso04_{tag}$) and transverse displacement from the primary vertex (d_{xy}). The standard tracker isolation is defined as the sum p_T of all tracks within $\Delta R < 0.4$ around the object relative to object p_T , which in this case is the tag muon. However, since the J/ψ produces collimated muons, the tracker isolation was modified as in Equation 5.7 to exclude the p_T contribution from the probe muon while calculating the tracker isolation of the tag muon. While the analysis used the PF isolation for event selection, the triggers typically use tracker isolation to define a cut during the trigger path. Therefore, the tracker isolation was preferred during these studies.

$$relIso04_{tag} = \frac{\sum_{trk}^{\Delta R < 0.4} p_{T,trk} - p_{T,tag} - p_{T,probe}}{p_{T,tag}} \quad (5.7)$$

Category	Condition
Tag muon	<ul style="list-style-type: none"> $p_T > 8 \text{ GeV}$, $\eta < 2.4$, Tight ID Matched to HLT_Mu8 or HLT_Mu17 or HLT_Mu20 $\text{relIso04} < 1$ $d_{xy} < 0.005$
Probe muon	$p_T > 5 \text{ GeV}$, $ \eta < 2.4$
Pair	<ul style="list-style-type: none"> $2.9 \text{ GeV} < m_{\mu\mu} < 3.3 \text{ GeV}$ $\Delta R \geq 0.04$
Fits	<ul style="list-style-type: none"> Signal: Crystal-Ball Background: Exponential

Table 5.15: Summary of the selection criteria used to identify muon pairs originating from J/ψ decays.

The second cut introduced is on the d_{xy} of the tag muon. This cut was required since data contains both prompt and non-prompt J/ψ while this study had access to prompt J/ψ simulation only. Furthermore, prompt J/ψ can be produced at high p_T , which is less likely to happen for non-prompt J/ψ , typically produced in B-hadron decays.

In comparison to the tag muon, the probe muon is required to only pass loose cuts such as the $p_T > 5 \text{ GeV}$ and $|\eta| < 2.4$. Since the muons in our analysis are also required to pass the loose muon ID and loose PF isolation cuts, the efficiency measurements are performed using probe muons that meet similar criteria. The selection also imposes a mass window for the tag-and-probe pair, centered on the known J/ψ mass. The mass distribution is then individually built for a range of ΔR bins from 0.04 to 0.4, for the passing and failing probes with $p_T > 26$ or 29 GeV, depending on the isolated single muon trigger threshold. Once the bins are filled, a signal+background fit is performed with exponential function for background and the Crystal-Ball function for the signal⁵. An example fit for simulation and data can be seen in Fig. 5.47. Since the number of high- p_T J/ψ being produced is small, the small ΔR bins are limited by statistics, which leads

⁵The Crystal-Ball function combines a Gaussian core with a power-law tail. The transition point is determined by α . The function is defined as $f(x; \alpha, n, \lambda) = N \exp(-\lambda^2/2)$ for $\lambda > -\alpha$ and $A(B - \lambda)^{-n}$ for $\lambda \leq -\alpha$ where $\lambda = \frac{x - \bar{x}}{\sigma}$, \bar{x} and σ are the mean and width of the Gaussian, and n determines the power of the power-law tail.

to bad fits in some cases. In those cases, a cut-and-count approach is instead taken to measure the efficiency as a function of the ΔR .

```

Fit status pass: 0, fail: 0
 $\chi^2/\text{ndof}$  pass: 1.965, fail: 0.734
KS pass: 0.829, fail: 0.747
eff = 0.5048 ± 0.0060
--- parameters
pass
- alphaP = 1.382 ± 0.872
- mP = 3.098 ± 0.001
- nBkgP = 270.333 ± 41.422
- nP = 2.984 ± 1.541
- nSigP = 3518.271 ± 70.437
- sP = 0.037 ± 0.001
fail
- alphaF = 0.828 ± 0.927
- mF = 3.101 ± 0.001
- nBkgF = 282.487 ± 45.574
- nF = 1.390 ± 0.463
- nSigF = 3451.138 ± 72.407
- sF = 0.036 ± 0.001

Fit status pass: 70, fail: 140
 $\chi^2/\text{ndof}$  pass: 0.977, fail: 0.522
KS pass: 0.482, fail: 0.951
eff = 0.5089 ± 0.0446
--- parameters
pass
- alphaP = 0.043 ± 1.001
- mP = 3.084 ± 0.005
- nBkgP = 94.356 ± 14.428
- nP = 0.979 ± 2.367
- nSigP = 73.679 ± 13.703
- sP = 0.030 ± 0.005
fail
- alphaF = -0.657 ± 2.991
- mF = 3.093 ± 0.005
- nBkgF = 31.908 ± 17.522
- nF = 0.161 ± 0.678
- nSigF = 71.096 ± 18.622
- sF = 0.032 ± 0.004

```

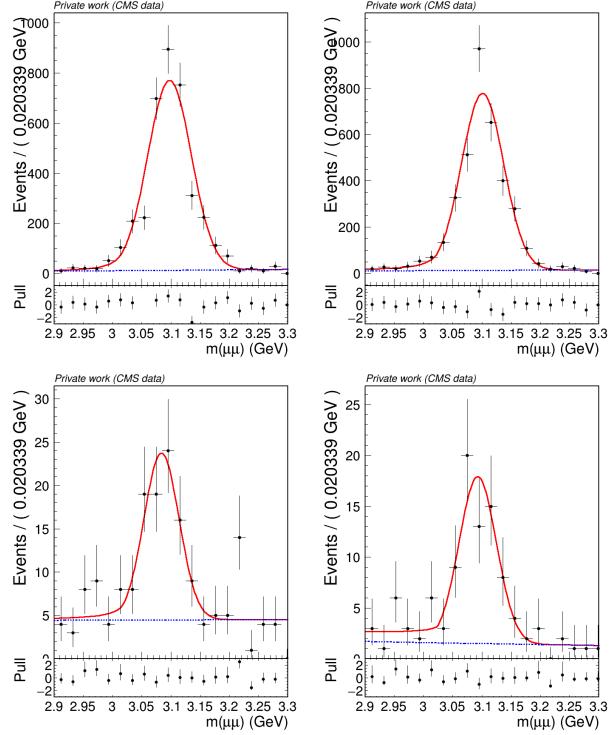


Figure 5.47: 2017 mass fits for simulation (top) and data (bottom) for passing events (left) and failing events (right) respective in the following ΔR bin: $0.04 < \Delta R < 0.07$, probe muon $p_T > 29 \text{ GeV}$ $|\eta| < 2.4$.

The results are shown in Fig. 5.48 for the four eras of Run 2. The trigger data-to-simulation scale factors for the bins where both data and simulation events are present, have been used as corrections in the analysis according to the $\Delta R_{\mu\mu}$ of the event. Additionally, for late 2016 era, the effective luminosity is very small ($L_{eff} \sim 0.06 fb^{-1}$) and leads to very little statistics. As the simulation efficiency in the early 2016 and late 2016 era is similar, the early 2016 results are applied for the full 2016 dataset.

As the study is limited by statistics, the total uncertainty on these measurements is calculated by varying the corrections by $\pm \sigma_{\text{stat.}}$ for each bin and esti-

mating the shift in the signal yield. This is derived for each era and each scalar mass considered. Since 2017 has the most effective luminosity, the statistically driven uncertainty is observed to be the smallest there ($\sim 4 - 5\%$) and goes up to $10 - 13\%$ for 2018.

Higgs boson p_T shape uncertainty and QCD scale variations

The Higgs boson p_T reweighting uncertainty is conventionally estimated by comparing the distribution after shifting the renormalization and factorization scales, μ_R and μ_F , by a factor of 0.5 and 2 and measuring the change in the signal yield. Fig. 5.49 shows the effect of different scale variations on the nominal Higgs boson p_T distribution. Note that the figure shown here includes contributions from the shape as well as normalization differences due to the scale variations. Using the conservative variations of $(\mu_R, \mu_F) = (0.5, 0.5)$ and $(\mu_R, \mu_F) = (2.0, 2.0)$, the signal yield varies by 17%. This is assigned as the systematic uncertainty. Since this variation contains the uncertainty of both the Higgs boson p_T shape and the normalization uncertainty on the cross-section, the theoretical uncertainty on the signal cross-section is not considered in addition to this.

Sample statistics

The limited sample size of the signal simulation introduces a statistical uncertainty to the analysis. This error is taken as the root of the squared sum of weights of the events that pass the analysis cuts.

For a quick overview, the table summarizing all the uncertainties is shown in Table 5.16.

5.8 Results

In the Section 5.6, the distributions of various quantities for the signal simulation, background prediction, and observed data provided a visual comparison of the agreement between the background-only expectation and the observed data. A statistical interpretation allows for the quantification of this comparison between the observed data and the background. This interpretation is vital to measure the

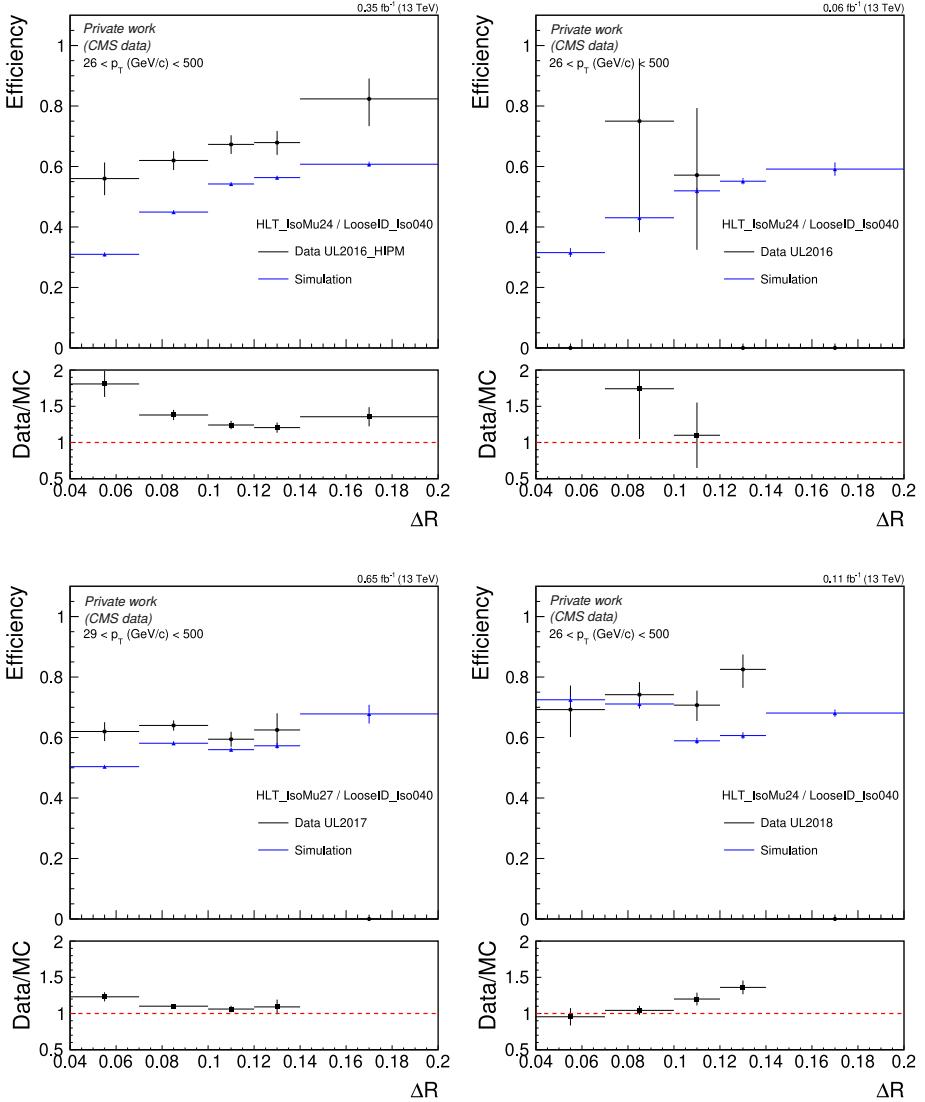


Figure 5.48: Trigger efficiency as a function of $\Delta R_{\mu\mu}$ for simulation (blue) and data (black) in the early 2016 (top left), late 2016 (top right), 2017 (bottom left) and 2018 (bottom right) eras of Run 2. The low effective luminosity in data leads to zero statistics in a few bins.

extent to which the analysis probes the parameter space of BSM scalar bosons produced from Higgs boson decays. In particular, upper limits on the branching

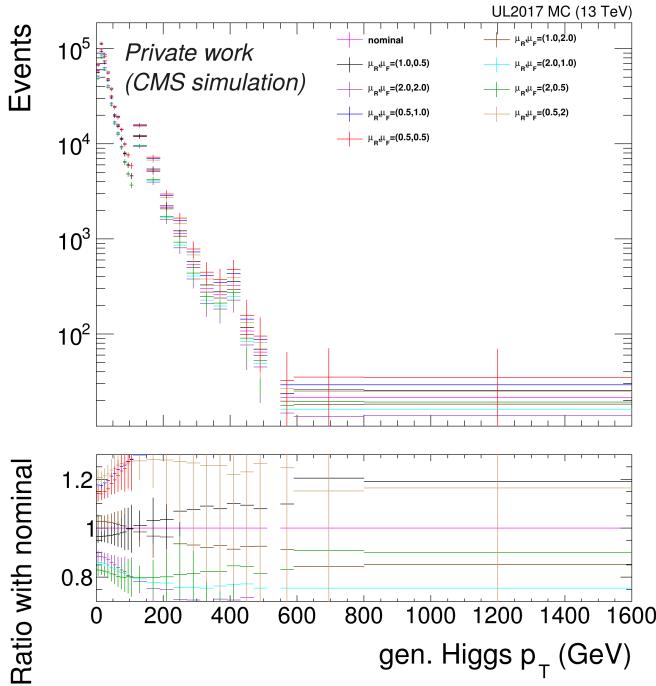


Figure 5.49: The generator level Higgs boson p_T of a reference signal sample with 9 independent variations of the factorization and renormalization scales (μ_R, μ_F) is shown with respect to the nominal weights.

fraction of the Higgs boson decay to the scalar bosons are evaluated using the signal yield, background prediction and the observed data. This is performed by the CMS `combine` tool [154]. In the following subsections, the details of the statistical procedure for limit setting are described, followed by the results of the analysis.

5.8.1 Limit setting procedure

The expected background and the potential signal contribution (according to the benchmark physics model taken) indicates that any event yield can be a very rare occurrence in comparison to the large number of pp collision events that have been recorded with the CMS experiment. This motivates using a Poisson distribution to construct a probability model for the expected/observed yield, denoted

Systematic uncertainty	Value	Comments
luminosity	1-3%	event-level
PU reweight	1%	event-level
Muon ID	0.6%	per muon
Reco. close muons	1%	event-level
Reco. displaced muons	5%	per muon
Trigger	5-13%	event-level
Reco. hadronic tracks	4-5%	per hadron
SV reco.	5-40%	per SV, based on category
Reco. scalar mass reso.	1-10%	per scalar mass, based on category
Higgs p_T shape	17%	event-level
Muon Isolation	2-12%	event-level, based on scalar mass
Hadron Isolation	6-16%	event-level, based on scalar mass
MC statistics	2-15%	event-level
Background statistics	15-100%	event-level, based on scalar mass & category
Transfer factor f_1	15-100%	only non- <i>prompt</i> categories
Transfer factor f_2	10-30%	only non- <i>prompt</i> categories

Table 5.16: Summary of systematic uncertainties.

by $p(x|n)$ where x is the observed event yield and n is the model-dependent expected yield. The background processes consisting of SM processes, denoted by b , and the signal model-dependent yield, denoted by s , combine to give the model-dependent expected yield such that $n = \mu s + b$. μ (also referred as r) is the signal strength modifier which for a background-only hypothesis would be $\mu = 0$. For the physics process considered here, the signal strength modifier is the only parameter of interest in the model and is a function of the branching ratio of the SM Higgs boson to the BSM light scalar bosons ($\mathcal{B}(H \rightarrow SS)$). By estimating the upper limits on the signal strength modifier, an inference on the $\mathcal{B}(H \rightarrow SS)$ can be made. Besides this, the model consists of other parameters, referred to as the nuisance parameters (v), which include all the systematic effects and corrections presented in the previous section, which modify the expected signal/background yield i.e. $b = b(v), s = s(v)$. In summary, the probability model considered is $Pois(x|\mu s(v) + b(v))$.

The next step is to define a function which measures how well the model describes the observed data. This is done by using the likelihood function defined as $\mathcal{L}(\mu, v) = Pois(x|\mu s(v) + b(v))$ where maximizing the likelihood would produce the Maximum Likelihood Estimators (MLE) for the signal strength and nuisance parameters (denoted as $\hat{\mu}$ and \hat{v}). Conversely, a negative log likelihood

function (NLL=-ln(\mathcal{L})) minimization would lead to the values of $\hat{\mu}$ and \hat{v} as well. The NLL function is computationally simpler and numerically more stable as it splits a product of likelihood functions into a sum of log likelihood functions. For instance, in this analysis, there are several categories considered and the combined likelihood would be of the form, $\mathcal{L}(\mu, v) \equiv \prod_{i=1}^{N_{\text{channels}}} \text{Pois}(x_i | \mu s_i(v) + b_i(v))$. The `combine` tool performs the NLL minimization to obtain the maximum likelihood estimators for the parameters of interest and the nuisance parameters.

The nuisance parameters in the model pose a challenge in the likelihood-based inference since they act as sources of systematic uncertainty. This would smear the estimate for the parameter of interest, thereby decreasing the sensitivity of the analysis. Depending on the type of nuisance parameter, each nuisance parameter can be approximately modeled with a probability distribution such that the likelihood is then modified as $\mathcal{L}(\mu, v) \equiv \prod_{i=1}^{N_{\text{channels}}} \text{Pois}(x_i | \mu, v) \prod_j p(y_j | v_j)$ where y is an auxiliary observable. There are several approaches to remove the nuisance parameters from the model ranging from a bayesian approach of integrating out the nuisance parameter from the likelihood function to a frequentist approach of performing a profile likelihood which involves measuring the MLE of the nuisance parameter for a fixed value of μ . The profile likelihood method in the `combine` tool has been used to deal with the nuisance parameters in this analysis.

The profile likelihood gives rise to $\hat{v}(\mu)$ which can also be interpreted as the MLE of v for a particular μ . Using this, a test-statistic is defined as follows:

$$q(\mu) = -2 \ln \frac{\mathcal{L}(x | \mu, \hat{v}(\mu))}{\mathcal{L}(x | \hat{\mu}, \hat{v})} \quad (5.8)$$

The test-statistic is then used for hypothesis testing and setting upper limits on the signal strength, μ . In particular, these limits are set at 95% confidence level using the CL_s method [158]. The procedure evaluates the p-values for the signal+background hypothesis and for the background-only hypothesis. This is calculated as $p_{s+b} = \int_{q_{\mu}^{\text{obs}}}^{\infty} f(q_{\mu} | \mu, v) dq_{\mu}$ where $f(q_{\mu} | \mu, v)$ is the probability distribution function for the test statistic q_{μ} . According to the convention used for `combine`, the p-value for the background is written as $1 - p_b =$

$\int_{q_\mu^{obs}}^{\infty} f(q_\mu | 0, v) dq_\mu$. The CL_s is then measured as the ratio of these two p-values: $CL_s = p_{s+b}/(1 - p_b)$. This is used to determine the 95% confidence interval for a signal hypothesis with signal strength μ . If the CL_s is found to be smaller than $\alpha = (1 - 0.95) = 0.05$, then the data is incompatible with the signal+background model and excludes the chosen μ with a 95% confidence level.

To understand and visualize the CL_s criterion and how it ties to testing the various signal models considered in the analysis, the model with BSM scalar boson of $m_S = 1.2$ GeV and $c\tau = 10$ mm is chosen as reference. For this model, the signal strength for the model is chosen such that the $\mathcal{B}(H \rightarrow SS) \sim 0.1\%$ and $\mathcal{B}(S \rightarrow \mu\mu), \mathcal{B}(S \rightarrow hh)$ according to theory [17].

The next step is the generation of the probability distribution functions, $f(q_\mu | \mu, v)$ and $f(q_\mu | 0, v)$, to compute the corresponding p-values. This is obtained by generating toy datasets which are MC pseudo-experiments, sampled from the probability distribution function of the signal+background model and the background only model. In this case, 20000 toys are generated for both the models and the q_μ for each of the toys is calculated.

The resulting distribution is shown in Fig. 5.50 (left) where the blue and red histograms show the q_μ distribution for signal+background and background only models, respectively. The black line in Fig. 5.50 (left) is the observed test-statistic for the full Run 2 dataset after the combination of all event categories.

The integral of the right-sided tail for the signal+background hypothesis, the p_{s+b} , is equal to 0.0309, indicating that the signal+background model is incompatible with the observed value. For the background only hypothesis, the integral of the right-sided tail in the distribution, $1 - p_b$, is found to be 0.7204, indicating that the data is compatible with the background-only hypothesis, meaning that any statistical fluctuations in the background are consistent with what is observed. The CL_s is found to be 0.0429, and allows for the exclusion of the considered signal model.

This is repeated for multiple values of $\mathcal{B}(H \rightarrow SS)$ and the corresponding signal strength, μ , to obtain the CL_s as a function of μ , as shown in Fig. 5.50 (right). It is seen that the CL_s decreases with increasing signal strength. It can therefore be said that all models with signal strength greater than $\mu_{@CL_s=0.05}$ are excluded. The central vertical red line points to the value of $\mu_{@CL_s=0.05}$. This

ultimately quantifies the sensitivity of the analysis and allows to set the upper limits on $\mathcal{B}(H \rightarrow SS)$ at 95% confidence level.

The full procedure listed here is repeated for the median (or 50% quantile), $\pm 1\sigma$ (or 16% quantile and 84% quantile) and $\pm 2\sigma$ (or 2.75% quantile and 97.5% quantile) values of the expected signal strength for comparison with the observed limits. This is then derived for each of the BSM scalar signal models considered in the analysis and the results for this are discussed in the next subsection.

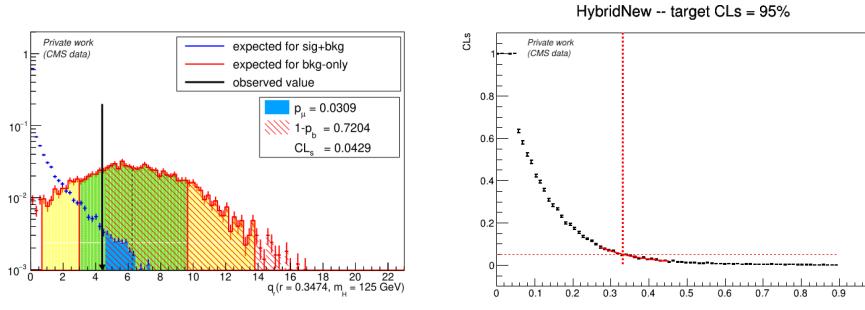


Figure 5.50: Left: Distribution of the test statistic for the full Run 2 data set and the signal model with BSM scalar boson of $m_S = 1.2$ GeV and $c\tau = 10$ mm for a given signal strength (μ or $r=0.3474$). The black arrow corresponds to the observed value of the test statistic while the blue and red distributions are $f(q_\mu|\mu, v)$ and $f(q_\mu|0, v)$ respectively, generated from toy datasets. Right: CL_s as a function of the signal strength for the observed data in full Run 2 for the signal model with BSM scalar boson of $m_S = 1.2$ GeV and $c\tau = 10$ mm. The vertical red lines indicate the observed limits and its statistical uncertainty coming from the number of toy datasets.

5.8.2 Expected and Observed Limits

The method described above, referred to as the ‘HybridNew’ method in the `combine` tool, is used to evaluate the limits for the full Run 2 dataset. The relevant systematic uncertainties, signal yield and estimated background count are given as input to the tool which returns the $\mu_{@CL_s=0.05}$. Using the relation $\mu = \mathcal{B}(H \rightarrow SS)(\text{obs})/\mathcal{B}(H \rightarrow SS)(\text{th.})$, the 95% confidence-level upper limits on $\mathcal{B}(H \rightarrow SS)$ are obtained as a function of the $c\tau$ for individual BSM scalar boson masses. Here, $\mathcal{B}(H \rightarrow SS)(\text{th.})$ corresponds to a reference value assigned

to the signal hypothesis. In the figures in Section 5.6, this was chosen to be 1% for illustration purposes. During the statistical interpretation, this is chosen according to the lifetime of the BSM scalar boson. Additionally, the branching fractions $\mathcal{B}(S \rightarrow \mu^+\mu^-)$, $\mathcal{B}(S \rightarrow hh)$ are taken from theory [17]. The contributions from Higgs boson production via vector boson fusion and in association with a vector boson are evaluated and accounted for as well, and increase the signal yields (across different m_S and lifetimes) by about 9.1%.

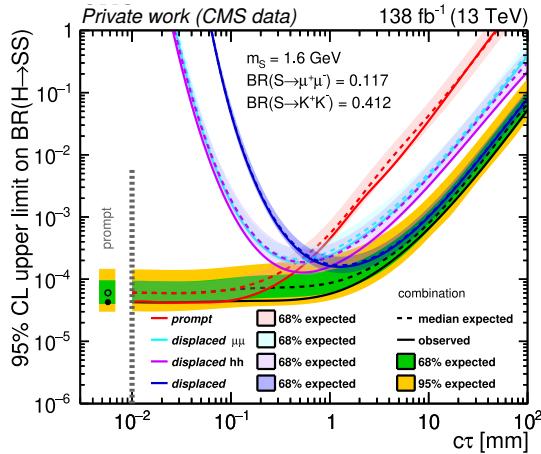


Figure 5.51: Observed (solid black lines) and expected (dashed black lines) exclusion limits on the branching fraction $\mathcal{B}(H \rightarrow SS)$ as function of signal proper lifetime $c\tau$ for $m_S = 1.6$ GeV for the most minimal extension of the SM Higgs sector [16], assuming $\mathcal{B}(S \rightarrow \mu^+\mu^-)$ and $\mathcal{B}(S \rightarrow K^+K^-)$ from Ref. [17]. The inner green (outer yellow) indicates the region containing 68% (95%) of the expected limits. The limits are obtained using the combination of all eras of Run 2 and the four different categories. The contribution of each category is shown with the expected (dashed) and observed (solid) lines along with the respective 68% uncertainty band.

As an example, Fig. 5.51 shows the combination limits (black) for the BSM scalar boson mass of 1.6 GeV as a function of $c\tau$ for the full Run 2 dataset. Lifetime reweighting (see Section 5.2) and interpolation between the points allows for a smooth curve, spanning from $c\tau = 0.01$ mm to $c\tau = 100$ mm. The combination for the limits in black is performed across the 4 event categories, each with its unique range of sensitivity as described below.

Starting with the *prompt* category (red), it is seen that the sensitivity drops quickly as the signal lifetime increases. This is expected since the category captures prompt-like events only. Below $c\tau \leq 0.1\text{mm}$, the limits start to plateau with respect to the lifetime and does not improve any further. This is related to travel distance of the BSM scalar bosons and the vertexing resolution of the detector. Given that the vertex resolution is somewhere between 10–20 μm , a conservative distance for a vertex to be reconstructed as ‘displaced’ and away from the interaction point, can be taken as 5 times the vertexing resolution and is equal to 0.1 mm. Following the relation $L = \beta\gamma c\tau$, and since $\beta\gamma \gg 1$ for the light scalar particles, it is not surprising that the limits plateau at some $c\tau$ between 0.01 and 0.1 mm. In this plateau region, the decay vertex is not reconstructed as a ‘displaced’ vertex anymore. Therefore, below 0.01 mm, the limits can be safely assumed to be constant and accordingly, the limits at 0.01 mm are interpreted as the sensitivity for a purely prompt signal.

The next two categories, *displaced* $\mu\mu$ (cyan) and *displaced* hh (magenta), provide similar sensitivity to the $0.1 \leq c\tau \leq 1\text{ mm}$ range. The combination of the two help to extend the sensitivity to this $c\tau$ range and make it comparable to what is observed for the *prompt*-like signals (i.e. $c\tau \leq 0.1\text{ mm}$). The *displaced* (blue) category increases the sensitivity to larger lifetimes and is the primary driver for the sensitivity to $c\tau \geq 1\text{ mm}$. For each of the non-*prompt* categories, a parabola-like shape is observed. This is because as the lifetime increases, more events start to populate these categories. However, once an optimal $c\tau$ range is reached, the sensitivity starts to drop as a larger number of events start to have one or more BSM scalar boson decays occurring outside the detector volume.

Figure 5.52 shows the final limits for BSM scalar bosons with $1.1 \leq m_S \leq 2\text{ GeV}$, where the final state hadrons are assumed to be charged kaons. Figure 5.53 shows the limits for BSM scalar bosons with $0.4 \leq m_S \leq 1\text{ GeV}$ with the charged pion assumption for the final state hadrons. It is observed that a similar trend is present for each BSM scalar boson mass – the best sensitivity is achieved at low lifetimes due to efficient reconstruction and triggering capabilities followed by worse limits at larger lifetimes where the vertices are more displaced and are harder to trigger on. In all cases, it is seen that the observed limits are in agreement with the expected limits within the 68% expected band.

To study the effect of the limits as a function of the BSM scalar boson mass, the observed limits are plotted in a two-dimensional map between the BSM scalar boson mass and the lifetime and are shown in Fig. 5.54. The black line connects the lifetimes where $\mathcal{B}(H \rightarrow SS) = 1\%$. Using this as an indicator, for most mass hypotheses, a limit on $\mathcal{B}(H \rightarrow SS)$ of the order of 10^{-3} (10^{-4}) for proper lifetimes up to $c\tau \sim 10$ mm ($c\tau \sim 1$ mm) is achieved. A drop in sensitivity is observed near 1 GeV due to a small $\mathcal{B}(S \rightarrow \mu\mu)$.

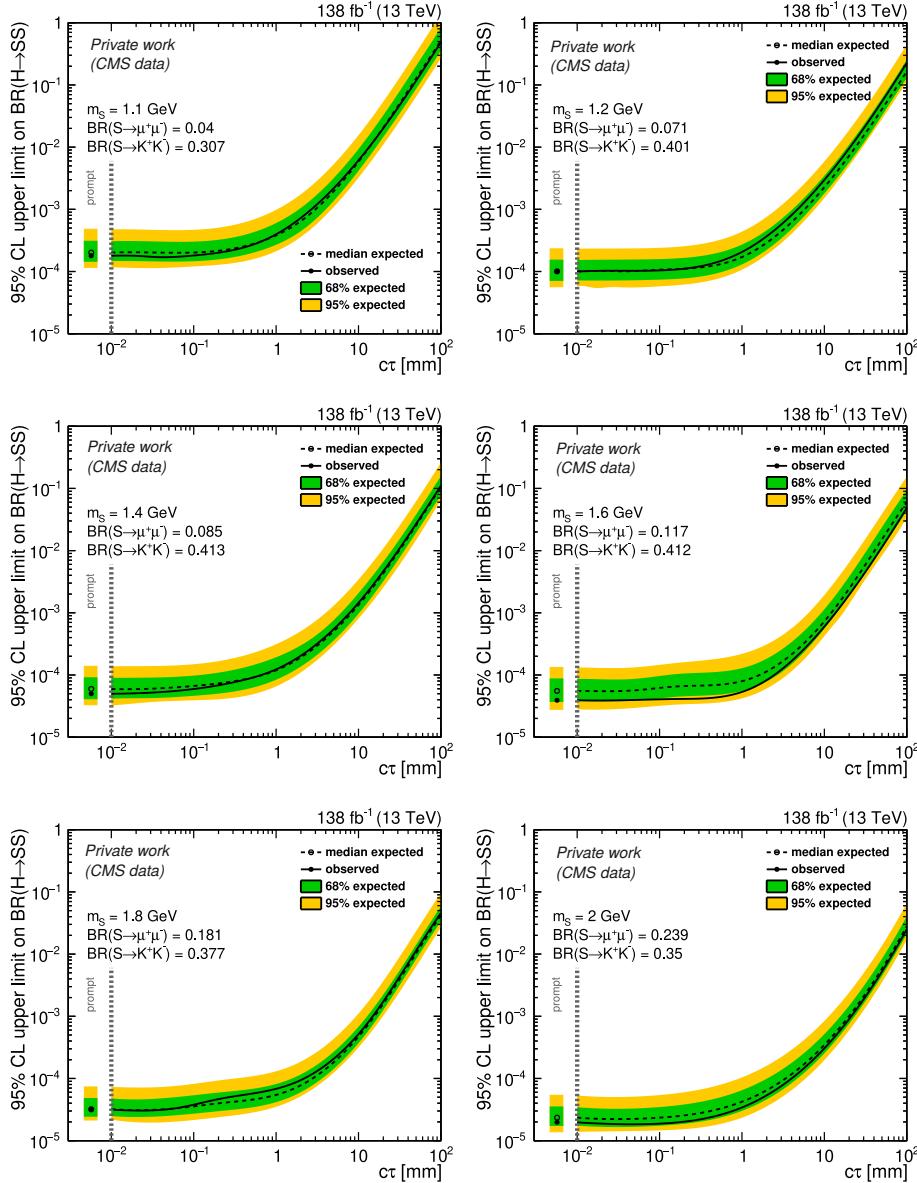


Figure 5.52: Observed (solid black line) and expected (dashed black line) exclusion limits on the branching fraction $\mathcal{B}(H \rightarrow SS)$ as function of signal proper lifetime $c\tau$ for $m_S = 1.1 - 2 \text{ GeV}$ for the most minimal extension of the SM Higgs sector [16], assuming $\mathcal{B}(S \rightarrow \mu^+ \mu^-)$ and $\mathcal{B}(S \rightarrow K^+ K^-)$ from Ref. [17]. The inner green (outer yellow) indicates the region containing 68% (95%) of the limits. The limits are obtained using the combination of all eras of Run 2 and the four different lifetime categories.

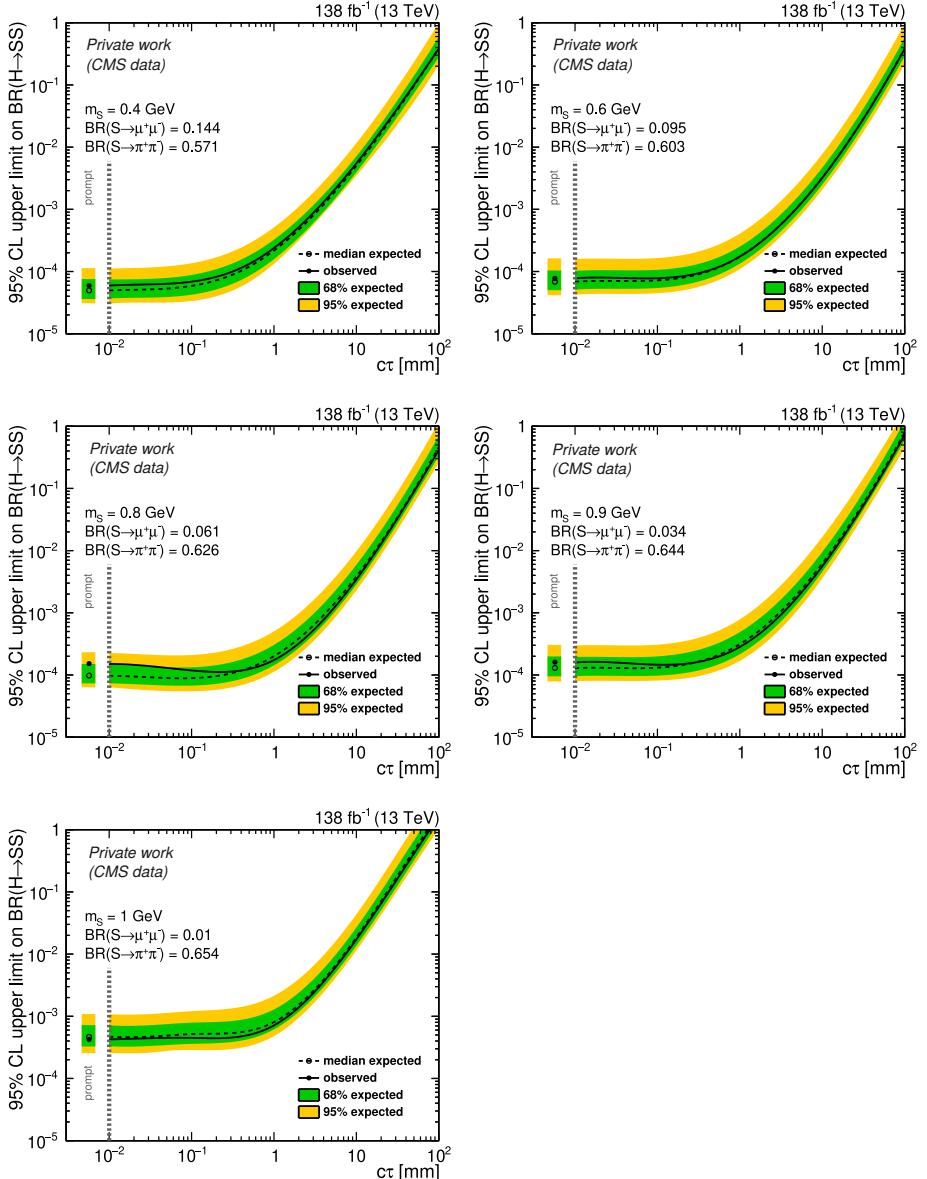


Figure 5.53: Observed (solid black line) and expected (dashed black line) exclusion limits on the branching fraction $\mathcal{B}(H \rightarrow SS)$ as function of signal proper lifetime $c\tau$ for $m_S = 0.4 - 1$ GeV for the most minimal extension of the SM Higgs sector [16], assuming $\mathcal{B}(S \rightarrow \mu^+\mu^-)$ and $\mathcal{B}(S \rightarrow \pi^+\pi^-)$ from Ref. [17]. The inner green (outer yellow) indicates the region containing 68% (95%) of the limits. The limits are obtained using the combination of all eras of Run 2 and the four different lifetime categories.

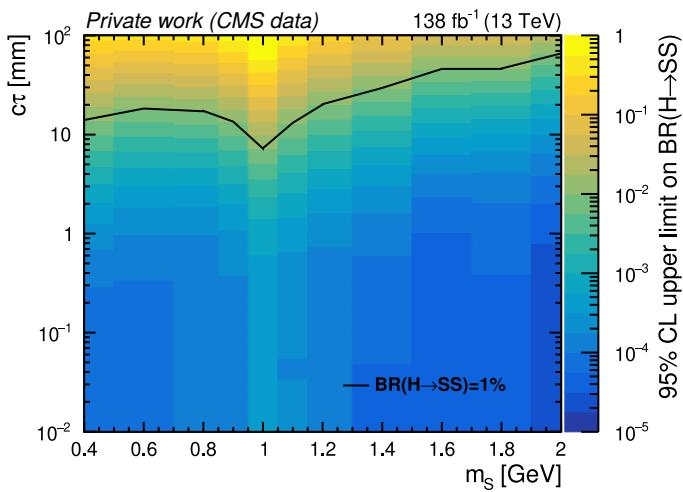


Figure 5.54: Observed upper limits on the branching fraction $\mathcal{B}(H \rightarrow SS)$ as function of signal mass and proper lifetime $c\tau$ for the most minimal extension of the SM Higgs sector [16], assuming $\mathcal{B}(S \rightarrow \mu^+\mu^-)$, $\mathcal{B}(S \rightarrow K^+K^-)$ and $\mathcal{B}(S \rightarrow \pi^+\pi^-)$ from Ref. [17]. The area under the solid black line denotes the region where the limits are smaller than 1%.

Conclusion and Outlook

This thesis presents two analyses carried out within the CMS experiment at the CERN Large Hadron Collider. The first one contributes to the calibration of jets at the CMS experiment by studying the jet energy resolution using a data-driven approach. The second analysis is a search for light scalar bosons, motivated by BSM scenarios where the Higgs boson decays exotically into a pair of such particles. The search targets events with a final state consisting of two muons and two hadrons, corresponding to the decays of the two light scalar bosons.

In the first analysis, a deep dive into the final stage of jet calibration, the jet energy resolution, was performed using the ultra-legacy (UL) 2018 data from the Run 2 data-taking period of the LHC. The jet energy resolution and its scale factor were studied with a focus on exploring the $p_T^{\text{jet}} < 100$ GeV region using Z+jet events. The first method, inspired from the p_T balance approach and developed to have a reduced bias towards pileup, provided scale factors of $\sim 1.1 - 1.3$. With the new detector calibrations in the ultra-legacy datasets, it was observed that the jet energy resolution showed a p_T dependence across all $|\eta|$. With the p_T balance approach, the results indicate larger scale factors as the p_T decreases. In parallel, a novel method based on the MET projection fraction (MPF) approach, currently in development within the CMS Collaboration, suggested the opposite trend of decreasing scale factors as the p_T decreases. The possibility of the difference in event selection was explored by applying the MPF method to the Z+jet event topology selected in the p_T balance study. The results were found to be compatible and indicated a JER scale factor of $\sim 1 - 1.1$.

To understand the impact of these two results and to check which of the results are possibly more accurate, a validation study using the leading jet p_T spectra in Z+jets was performed by applying both the JER scale factors to simulation and comparing the results. While the method based on the MPF approach shows a better agreement with data, the difference was found to be small and the uncertainty on this is covered within other uncertainties on the simulation. This leads to the conclusion that large-scale observables such as the jet p_T spectra are expected to show small impact with the application of either of the two JER results. To further understand the impact of the two results, a similar study using a more resolution-sensitive variable such as the reconstructed width of a hadronic resonance like the W boson decaying to two jets, would provide more insight.

In the second analysis, a search for the Higgs-mediated production of a pair of light scalar bosons with masses in the range of 0.4–2 GeV was performed using Run 2 data collected by the CMS experiment. In this mass range, the analysis accesses a unique hadronic decay mode where the low mass of the scalar restricts hadronization and dominantly allows decays to only a pair of light hadrons. However, since the pp collision produces a dense hadronic environment, it is challenging to select events at the trigger level with fully hadronic final-states. Instead, the subdominant channel to muons is targeted for one of the scalar bosons while the decay to charged hadrons is considered for the other scalar boson. Furthermore, since the coupling to SM particles depends on the mixing angle θ , which eventually determines the lifetime of the particle, the search is performed for scalar bosons up to a proper lifetime of $c\tau = 100$ mm, targeting events which decay within the CMS tracker volume.

An event selection based on the boosted topology, the isolation, and mass constraints between the reconstructed dimuon and dihadron resonances was devised to reduce the considerable background coming largely from QCD multijet processes, reducing it to nearly zero. A statistical analysis of the signal prediction, background prediction and the observed number of events is performed to set upper limits on the branching fraction of the Higgs boson decaying to light scalar bosons ($\mathcal{B}(H \rightarrow SS)$), considering the benchmark signal model in Ref. [17].

This search excludes at 95% confidence level branching fractions of the

Higgs boson to light scalars greater than 10^{-4} and covers a largely unexplored phase space of light particles, providing complementarity to several Run 2 searches on long-lived particles (BSM particles that travel a certain distance before decaying) (see Fig. 5.55). In particular, this analysis provides complementarity to the dimuon scouting results shown by probing all masses within the mass range of 0.4 to 2 GeV, some of which were masked in Ref. [62] and to the Hadronic MS analysis [61] which is sensitive to a larger proper lifetime.

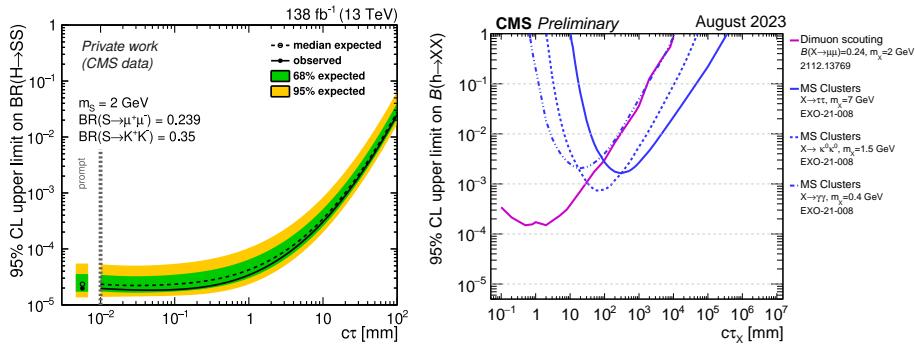


Figure 5.55: Left: Upper limits on the branching fraction of exotic Higgs decays to light scalar bosons as function of signal proper lifetime $c\tau$ for $m_S = 2$ GeV for the analysis presented in this thesis. Right: Upper limits on the branching fraction of exotic Higgs boson decays to long-lived particles as a function of the lifetime for a few selected mass points from various CMS analyses [61, 62]. A brief discussion has been presented in Section 1.2.3. Figure (right) taken from Ref. [63].

While this analysis provides a good sensitivity over various masses and lifetimes of the light scalar boson, there is still scope for improvements. One of the limitations of the analysis is the smaller sensitivity to larger lifetimes, as seen in Fig. 5.55 (left). The primary cause behind a lower sensitivity at $c\tau \geq 10$ mm is that the most optimal trigger for the analysis, the single muon trigger, is tuned towards promptly decaying signatures and rejects events where the muons are produced at a distance from the primary vertex. A large improvement on this limitation is expected in Run 3 with the addition of new displaced dimuon triggers [159], which are designed to select displaced muons and are also equipped with lower thresholds as compared to their counterparts in 2018 data-taking era.

Additionally, there are also alternate data streams such as data parking [91] which offer low p_T muon triggers ($p_T \gtrsim 7$ GeV), which can further improve the sensitivity for the search.

In the coming years, searches for light BSM particles or dark sector particles at CERN will be extensively pursued, driven by advances in detector technology and possibly from several proposed detectors at the LHC. Within the CMS experiment, the upcoming HL-LHC upgrades for the CMS tracker will enable particle track information from the outer regions of the tracker to be available already at the first level of trigger (L1) [160]. The ability to trigger directly on displaced tracks and/or displaced vertices will significantly boost the sensitivity for light long-lived particles in general. In parallel, several dedicated experiments such as CODEX-b [161], MATHUSLA [162], ANUBIS [163] and others have been proposed for long-lived particle searches. As these detectors would be placed away from the collision points, they can significantly extend the reach to lifetimes and correspondingly, very small mixing angles, that have not been searched for before. A comprehensive review of these future experiments can be found in Ref. [164]. Beyond collider experiments, many other experiments around the world are also probing the dark sector. Together, the upcoming developments and studies on understanding the nature of the dark sector will have the potential to unravel something new and pave the way beyond the Standard Model.

Author Contributions

The work presented in Chapter 4 on the jet energy resolution in the CMS experiment has been performed by me. Accordingly, all the figures and the methods presented in the Chapter have been performed by me. In Chapter 5, I describe the analysis that is presented in Physics Analysis Summary listed below. The figures and the results in the Chapter have been generated by me.

Papers

- CMS collaboration. “Search for light scalar particles from Higgs boson decays in exclusive final states with two muons and two hadrons”, *Physics Analysis Summary*, CMS-PAS-EXO-24-034: A preliminary version of the analysis performed in this thesis has been released. The corresponding paper is in preparation and will be submitted to publication.

Presentations

- “Recent results for long-lived and unconventional signatures in CMS”, *13th Edition of the Large Hadron Collider Physics Conference 2025*, 5-9 May, 2025
- “Displaced vertices+displaced tracks”, *CMS Exotica Workshop 2024*, 16-18 December, 2024

- “Exclusive search for light scalars via the Higgs portal decaying to pairs of muons and hadrons at the CMS experiment”, *Belgian Physical Society general scientific meeting 2024*, 29 May, 2024
- “JER pt-dependence: road to high precision jecs ”, *CMS JetMET workshop 2023*, 15-17 May, 2023

Bibliography

- [1] D. J. Griffiths, “Introduction to elementary particles”. Wiley, New York, NY, 2008.
- [2] J. J. Thomson, “Cathode rays”, *Philosophical Magazine* **44** (1897) 293–316, doi:10.1080/14786449708621070.
- [3] P. A. M. Dirac, “The quantum theory of the electron”, *Proc. Roy. Soc. Lond. A* **117** (1928) 610–624, doi:10.1098/rspa.1928.0023.
- [4] C. D. Anderson, “The positive electron”, *Phys. Rev.* **43** (1933) 491–494, doi:10.1103/PhysRev.43.491.
- [5] M. Gell-Mann, “A schematic model of baryons and mesons”, *Phys. Lett.* **8** (1964) 214–215, doi:10.1016/S0031-9163(64)92001-3.
- [6] G. Zweig, “An SU(3) model for strong interaction symmetry and its breaking. Version 2”, *CERN Report* (1964), no. 8182/TH.401,.
- [7] J. D. Bjorken and S. L. Glashow, “Elementary Particles and SU(4)”, *Phys. Lett.* **11** (1964) 255–257, doi:10.1016/0031-9163(64)90433-0.
- [8] E598 Collaboration Collaboration, “Experimental observation of a heavy particle J”, *Phys. Rev. Lett.* **33** (1974) 1404–1406, doi:10.1103/PhysRevLett.33.1404.

- [9] SPEAR Collaboration Collaboration, “Discovery of a narrow resonance in e^+e^- annihilation”, *Phys. Rev. Lett.* **33** (1974) 1406–1408,
doi:[10.1103/PhysRevLett.33.1406](https://doi.org/10.1103/PhysRevLett.33.1406).
- [10] CMS Collaboration, “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”, *Phys. Lett. B* **716** (2012) 30–61,
doi:[10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021),
arXiv:[1207.7235](https://arxiv.org/abs/1207.7235).
- [11] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Phys. Lett. B* **716** (2012) 1–29,
doi:[10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020),
arXiv:[1207.7214](https://arxiv.org/abs/1207.7214).
- [12] P. W. Higgs, “Broken symmetries, massless particles and gauge fields”, *Phys. Lett.* **12** (1964) 132–133,
doi:[10.1016/0031-9163\(64\)91136-9](https://doi.org/10.1016/0031-9163(64)91136-9).
- [13] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons”, *Phys. Rev. Lett.* **13** (1964) 508–509,
doi:[10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [14] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons”, *Phys. Rev. Lett.* **13** (1964) 321–323,
doi:[10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321).
- [15] F. Halzen and A. Martin, “Quarks & Leptons: An introductory course in modern particle physics”. John Wiley & Sons, New York, USA, 1984.
- [16] M. W. Winkler, “Decay and detection of a light scalar boson mixing with the Higgs boson”, *Phys. Rev. D* **99** (2019) 015018,
doi:[10.1103/PhysRevD.99.015018](https://doi.org/10.1103/PhysRevD.99.015018), arXiv:[1809.01876](https://arxiv.org/abs/1809.01876).
- [17] Y. Gershtein, S. Knapen, and D. Redigolo, “Probing naturally light singlets with a displaced vertex trigger”, *Phys. Lett. B* **823** (2021)

136758, doi:10.1016/j.physletb.2021.136758,
arXiv:2012.07864.

[18] I. Boiarska et al., “Phenomenology of GeV-scale scalar portal”, *JHEP* **11** (2019) 162, doi:10.1007/JHEP11(2019)162,
arXiv:1904.10447.

[19] I. Neutelings, “Standard model”, February, 2023. Accessed: 2025-10-25.
https://tikz.net/sm_particles/.

[20] S. Knapen and S. Lowette, “A Guide to Hunting Long-Lived Particles at the LHC”, *Ann. Rev. Nucl. Part. Sci.* **73** (2023) 421–449,
doi:10.1146/annurev-nucl-101920-013011,
arXiv:2212.03883.

[21] S. L. Glashow, “Partial Symmetries of Weak Interactions”, *Nucl. Phys.* **22** (1961) 579–588, doi:10.1016/0029-5582(61)90469-2.

[22] S. Weinberg, “A Model of Leptons”, *Phys. Rev. Lett.* **19** (1967) 1264–1266, doi:10.1103/PhysRevLett.19.1264.

[23] A. Salam, “Weak and Electromagnetic Interactions”, *Conf. Proc. C* **680519** (1968) 367–377, doi:10.1142/9789812795915_0034.

[24] J. Ellis, M. K. Gaillard, and D. V. Nanopoulos, “A Historical Profile of the Higgs Boson”, pp. 255–274. 2016. arXiv:1504.07217.
doi:10.1142/9789814733519_0014.

[25] M. E. Peskin and D. V. Schroeder, “An Introduction to quantum field theory”. Addison-Wesley, Reading, USA, 1995.
doi:10.1201/9780429503559.

[26] SNO Collaboration, “Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory”, *Phys. Rev. Lett.* **89** (2002) 011301,
doi:10.1103/PhysRevLett.89.011301,
arXiv:nucl-ex/0204008.

- [27] Super-Kamiokande Collaboration, “Evidence for oscillation of atmospheric neutrinos”, *Phys. Rev. Lett.* **81** (1998) 1562–1567, doi:10.1103/PhysRevLett.81.1562, arXiv:hep-ex/9807003.
- [28] D. Brout and et. al., “The pantheon+ analysis: Cosmological constraints”, *The Astrophysical Journal* **938** (oct, 2022) 110, doi:10.3847/1538-4357/ac8e04.
- [29] K. Freese, “Review of Observational Evidence for Dark Matter in the Universe and in upcoming searches for Dark Stars”, *EAS Publ. Ser.* **36** (2009) 113–126, doi:10.1051/eas/0936016, arXiv:0812.4005.
- [30] P. Salucci, “The distribution of dark matter in galaxies”, *Astron. Astrophys. Rev.* **27** (2019), no. 1, 2, doi:10.1007/s00159-018-0113-1, arXiv:1811.08843.
- [31] P. Fayet, “Supergauge Invariant Extension of the Higgs Mechanism and a Model for the electron and Its Neutrino”, *Nucl. Phys. B* **90** (1975) 104–124, doi:10.1016/0550-3213(75)90636-7.
- [32] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter”, *Phys. Rept.* **267** (1996) 195–373, doi:10.1016/0370-1573(95)00058-5, arXiv:hep-ph/9506380.
- [33] H. P. Nilles, “Supersymmetry, Supergravity and Particle Physics”, *Phys. Rept.* **110** (1984) 1–162, doi:10.1016/0370-1573(84)90008-5.
- [34] R. Barbieri, “Looking Beyond the Standard Model: The Supersymmetric Option”, *Riv. Nuovo Cim.* **11N4** (1988) 1–45, doi:10.1007/BF02725953.
- [35] H. E. Haber and G. L. Kane, “The Search for Supersymmetry: Probing Physics Beyond the Standard Model”, *Phys. Rept.* **117** (1985) 75–263, doi:10.1016/0370-1573(85)90051-1.

- [36] G. F. Giudice and A. Romanino, “Split supersymmetry”, *Nucl. Phys. B* **699** (2004) 65–89, doi:10.1016/j.nuclphysb.2004.08.001, arXiv:hep-ph/0406088. [Erratum: Nucl.Phys.B 706, 487–487 (2005)].
- [37] J. F. Gunion and H. E. Haber, “Higgs Bosons in supersymmetric models (III). Decays into neutralinos and charginos”, *Nucl. Phys. B* **307** (1988) 445, doi:10.1016/0550-3213(88)90259-3. [Erratum: Nucl.Phys.B 402, 569 (1993)].
- [38] CMS Collaboration, “Dark sector searches with the CMS experiment”, *Phys. Rept.* **1115** (2025) 448–569, doi:10.1016/j.physrep.2024.09.013, arXiv:2405.13778.
- [39] M. Shaposhnikov and I. Tkachev, “The vMSM, inflation, and dark matter”, *Phys. Lett. B* **639** (2006) 414, doi:10.1016/j.physletb.2006.06.063, arXiv:hep-ph/0604236.
- [40] F. Bezrukov and D. Gorbunov, “Light inflaton Hunter’s Guide”, *JHEP* **05** (2010) 010, doi:10.1007/JHEP05(2010)010, arXiv:0912.0390.
- [41] F. Bezrukov and D. Gorbunov, “Light inflaton after LHC8 and WMAP9 results”, *JHEP* **07** (2013) 140, doi:10.1007/JHEP07(2013)140, arXiv:1303.4395.
- [42] U. Ellwanger, C. Hugonie, and A. M. Teixeira, “The Next-to-Minimal Supersymmetric Standard Model”, *Phys. Rept.* **496** (2010) 1, doi:10.1016/j.physrep.2010.07.001, arXiv:0910.1785.
- [43] P. Fayet, “Supergauge Invariant Extension of the Higgs Mechanism and a Model for the Electron and Its Neutrino”, *Nucl. Phys. B* **90** (1975) 104, doi:10.1016/0550-3213(75)90636-7.

- [44] M. Pospelov, A. Ritz, and M. B. Voloshin, “Secluded WIMP Dark Matter”, *Phys. Lett. B* **662** (2008) 53,
doi:10.1016/j.physletb.2008.02.052,
arXiv:0711.4866.
- [45] G. Krnjaic, “Probing light thermal dark matter with a Higgs portal mediator”, *Phys. Rev. D* **94** (Oct, 2016) 073009,
doi:10.1103/PhysRevD.94.073009.
- [46] P. W. Graham, D. E. Kaplan, and S. Rajendran, “Cosmological relaxation of the electroweak scale”, *Phys. Rev. Lett.* **115** (Nov, 2015) 221801,
doi:10.1103/PhysRevLett.115.221801.
- [47] M. B. Voloshin, “Once gain About the Role of Gluonic Mechanism in Interaction of Light Higgs Boson with Hadrons”, *Sov. J. Nucl. Phys.* **44** (1986) 478.
- [48] S. Raby and G. B. West, “The Branching Ratio for a Light Higgs to Decay Into $\mu^+\mu^-$ Pairs”, *Phys. Rev. D* **38** (1988) 3488,
doi:10.1103/PhysRevD.38.3488.
- [49] A. Monin, A. Boyarsky, and O. Ruchayskiy, “Hadronic decays of a light Higgs-like scalar”, *Phys. Rev. D* **99** (2019), no. 1, 015019,
doi:10.1103/PhysRevD.99.015019, arXiv:1806.07759.
- [50] J. F. Donoghue, J. Gasser, and H. Leutwyler, “The Decay of a Light Higgs Boson”, *Nucl. Phys. B* **343** (1990) 341–368,
doi:10.1016/0550-3213(90)90474-R.
- [51] T. N. Truong and R. S. Willey, “Branching Ratios for Decays of Light Higgs Bosons”, *Phys. Rev. D* **40** (1989) 3635,
doi:10.1103/PhysRevD.40.3635.
- [52] J. F. Gunion, H. E. Haber, G. L. Kane, and S. Dawson, “The Higgs Hunter’s Guide”, volume 80. 2000.
doi:10.1201/9780429496448, ISBN 978-0-429-49644-8.

[53] P. N. Bhattacharya and J. D. Wells, “Depleted Higgs boson: Searches for universal coupling suppression, invisible decays, and mixed-in scalars”, *Phys. Rev. D* **107** (2023), no. 5, 055022, doi:10.1103/PhysRevD.107.055022, arXiv:2204.03435.

[54] T. Ferber, A. Grohsjean, and F. Kahlhoefer, “Dark Higgs bosons at colliders”, *Prog. Part. Nucl. Phys.* **136** (2024) 104105, doi:10.1016/j.ppnp.2024.104105, arXiv:2305.16169.

[55] LHCb Collaboration, “Search for hidden-sector bosons in $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ decays”, *Phys. Rev. Lett.* **115** (2015), no. 16, 161802, doi:10.1103/PhysRevLett.115.161802, arXiv:1508.04094.

[56] LHCb Collaboration, “Search for long-lived scalar particles in $B^+ \rightarrow K^+ \chi(\mu^+ \mu^-)$ decays”, *Phys. Rev. D* **95** (2017), no. 7, 071101, doi:10.1103/PhysRevD.95.071101, arXiv:1612.07818.

[57] Belle-II Collaboration, “Search for a long-lived spin-0 mediator in $b \rightarrow s$ transitions at the Belle II experiment”, *Phys. Rev. D* **108** (2023), no. 11, L111104, doi:10.1103/PhysRevD.108.L111104, arXiv:2306.02830.

[58] CMS Collaboration, “A portrait of the Higgs boson by the CMS experiment ten years after the discovery.”, *Nature* **607** (2022), no. 7917, 60–68, doi:10.1038/s41586-022-04892-x, arXiv:2207.00043. [Erratum: *Nature* 623, (2023)].

[59] CMS Collaboration, “Search for invisible decays of the Higgs boson produced via vector boson fusion in proton-proton collisions at $s=13$ TeV”, *Phys. Rev. D* **105** (2022), no. 9, 092007, doi:10.1103/PhysRevD.105.092007, arXiv:2201.11585.

[60] LHC Higgs Cross Section Working Group Collaboration, “Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector”, *CERN Yellow Rep. Monogr.* **2** (2017) 1–869, doi:10.23731/CYRM-2017-002, arXiv:1610.07922.

- [61] CMS Collaboration, “Search for long-lived particles decaying in the CMS muon detectors in proton-proton collisions at $\sqrt{s}=13$ TeV”, *Phys. Rev. D* **110** (2024) 032007, doi:10.1103/PhysRevD.110.032007, arXiv:2402.01898.
- [62] CMS Collaboration, “Search for long-lived particles decaying into muon pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV collected with a dedicated high-rate data stream”, *JHEP* **04** (2022) 062, doi:10.1007/JHEP04(2022)062, arXiv:2112.13769.
- [63] CMS Collaboration, “CMS Exotica Summary plots for 13 TeV data ”. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/SummaryPlotsEX013TeV>, 2025. Accessed: October 26, 2025.
- [64] CMS Collaboration, “Model-independent search for pair production of new bosons decaying into muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **12** (2024) 172, doi:10.1007/JHEP12(2024)172, arXiv:2407.20425.
- [65] L. Evans and P. Bryant, “LHC Machine”, *Journal of Instrumentation* **3** (aug, 2008) S08001, doi:10.1088/1748-0221/3/08/S08001.
- [66] R. Vogt, “Heavy ion physics at the LHC”, *Nucl. Phys. A* **752** (2005) 447–456, doi:10.1016/j.nuclphysa.2005.02.050, arXiv:hep-ph/0412301.
- [67] ALICE Collaboration, “ALICE: Technical Proposal”. LHC technical proposal. CERN, Geneva, 1995.
- [68] ATLAS Collaboration, “ATLAS: Technical Proposal”. LHC technical proposal. CERN, Geneva, 1994.
doi:10.17181/CERN.NR4P.BG9K.
- [69] CMS Collaboration, “CMS: Technical Proposal”. LHC technical proposal. CERN, Geneva, 1994.
- [70] LHCb Collaboration, “LHCb: Technical Proposal”. LHC technical proposal. CERN, Geneva, 1998.

[71] CERN, “The CERN Accelerator Complex”.
<https://home.cern/science/accelerators>, 2025.
Accessed: 2025-11-01.

[72] O. S. Brüning et al., “LHC Design Report”. CERN Yellow Reports: Monographs. CERN, Geneva, 2004.
doi:10.5170/CERN-2004-003-V-1.

[73] O. Brüning, M. Klein, S. Myers, and L. Rossi, “70 years at the high-energy frontier with the CERN accelerator complex”, *Nature Rev. Phys.* **6** (2024), no. 10, 628–637,
doi:10.1038/s42254-024-00758-5.

[74] CMS Collaboration, “Public CMS Luminosity Information”.
<https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>.

[75] CMS Collaboration, “The CMS Experiment at the CERN LHC”, *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.

[76] CMS Collaboration, “CMS”, technical report, CERN, Geneva, 1992.
Open presentation to the LHCC 5 November 1992, M. Della Negra/CERN, CMS Spokesman.

[77] CMS Collaboration, “Detector Drawings”, 2012. CMS Collection.

[78] I. Neutelings, “CMS coordinate system”, 2023. Accessed: 2025-10-25.
https://tikz.net/axis3d_cms/.

[79] CMS Collaboration, “The CMS tracker system project”. Technical design report. CMS. CERN, Geneva, 1997.

[80] CMS Tracker Group Collaboration, “The CMS Phase-1 Pixel Detector Upgrade”, *JINST* **16** (2021), no. 02, P02027,
doi:10.1088/1748-0221/16/02/P02027,
arXiv:2012.14304.

- [81] CMS Collaboration, “CMS DPG Tracker Public Results”.
<https://twiki.cern.ch/twiki/bin/view/CMSPublic/DPGResultsTRK>, 2025. Accessed: 2025-11-01.
- [82] CMS Collaboration, “Simulation of the Silicon Strip Tracker pre-amplifier in early 2016 data”, CMS Detector Performance Summary, 2020. CMS-DP-2020-045.
- [83] CMS Collaboration, “Operation and performance of the CMS silicon strip tracker with proton-proton collisions at the CERN LHC”, *JINST* **20** (2025), no. 08, P08027,
doi:10.1088/1748-0221/20/08/P08027,
arXiv:2506.17195.
- [84] CMS Collaboration, “The CMS electromagnetic calorimeter project”. Technical design report. CMS. CERN, Geneva, 1997.
- [85] CMS Collaboration, “Performance of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 13$ TeV”, *JINST* **19** (2024), no. 09, P09004, *doi*:10.1088/1748-0221/19/09/P09004,
arXiv:2403.15518.
- [86] CMS Collaboration, “The CMS hadron calorimeter project”. Technical design report. CMS. CERN, Geneva, 1997.
- [87] CMS Collaboration, “Calibration of the CMS hadron calorimeters using proton-proton collision data at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020), no. 05, P05002, *doi*:10.1088/1748-0221/15/05/P05002,
arXiv:1910.00079.
- [88] A. Colaleo, A. Safonov, A. Sharma, and M. Tytgat, “CMS Technical Design Report for the Muon Endcap GEM Upgrade”, technical report, 2015.
- [89] CMS Collaboration, “The CMS trigger system”, *JINST* **12** (2017), no. 01, P01020, *doi*:10.1088/1748-0221/12/01/P01020,
arXiv:1609.02366.

[90] H. Kwon, “Performance of the CMS Level-1 Trigger during Run 2”, *PoS ICHEP2020* (2021) 806, doi:10.22323/1.390.0806.

[91] CMS Collaboration, “Enriching the physics program of the CMS experiment via data scouting and data parking”, *Phys. Rept.* **1115** (2025) 678–772, doi:10.1016/j.physrep.2024.09.006, arXiv:2403.16134.

[92] A. Buckley et al., “General-purpose event generators for LHC physics”, *Phys. Rept.* **504** (2011) 145–233, doi:10.1016/j.physrep.2011.03.005, arXiv:1101.2599.

[93] R. K. Ellis, W. J. Stirling, and B. R. Webber, “QCD and collider physics”, volume 8. Cambridge University Press, 2, 2011. doi:10.1017/CBO9780511628788.

[94] P. Nason, “A New method for combining NLO QCD with shower Monte Carlo algorithms”, *JHEP* **11** (2004) 040, doi:10.1088/1126-6708/2004/11/040, arXiv:hep-ph/0409146.

[95] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with Parton Shower simulations: the POWHEG method”, *JHEP* **11** (2007) 070, doi:10.1088/1126-6708/2007/11/070, arXiv:0709.2092.

[96] S. Alioli, P. Nason, C. Oleari, and E. Re, “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”, *JHEP* **06** (2010) 043, doi:10.1007/JHEP06(2010)043, arXiv:1002.2581.

[97] J. Alwall et al., “Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions”, *Eur. Phys. J. C* **53** (2008) 473–500, doi:10.1140/epjc/s10052-007-0490-5, arXiv:0706.2569.

- [98] J. Alwall et al., “MadGraph 5 : Going Beyond”, *JHEP* **06** (2011) 128, doi:10.1007/JHEP06(2011)128, arXiv:1106.0522.
- [99] J. Alwall et al., “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, *JHEP* **07** (2014) 079, doi:10.1007/JHEP07(2014)079, arXiv:1405.0301.
- [100] R. Frederix and S. Frixione, “Merging meets matching in MC@NLO”, *JHEP* **12** (2012) 061, doi:10.1007/JHEP12(2012)061, arXiv:1209.6215.
- [101] NNPDF Collaboration, “Parton distributions from high-precision collider data”, *Eur. Phys. J. C* **77** (2017), no. 10, 663, doi:10.1140/epjc/s10052-017-5199-5, arXiv:1706.00428.
- [102] B. Andersson, “The Lund Model”, volume 7. Cambridge University Press, 1998. doi:10.1017/9781009401296.
- [103] T. Sjöstrand et al., “An introduction to PYTHIA 8.2”, *Comput. Phys. Commun.* **191** (2015) 159, doi:10.1016/j.cpc.2015.01.024, arXiv:1410.3012.
- [104] T. Sjöstrand and M. van Zijl, “A Multiple Interaction Model for the Event Structure in Hadron Collisions”, *Phys. Rev. D* **36** (1987) 2019, doi:10.1103/PhysRevD.36.2019.
- [105] S. Höche, “Introduction to parton-shower event generators”, in *Journeys Through the Precision Frontier: Amplitudes for Colliders*, pp. 235–295. 2015. arXiv:1411.4085. doi:10.1142/9789814678766_0005.
- [106] GEANT4 Collaboration, “GEANT4 - A Simulation Toolkit”, *Nucl. Instrum. Meth. A* **506** (2003) 250–303, doi:10.1016/S0168-9002(03)01368-8.

[107] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017), no. 10, P10003, doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965.

[108] CMS Collaboration, “Description and Performance of Track and Primary-Vertex Reconstruction with the CMS Tracker”, *JINST* **9** (2014), no. 10, P10009, doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569.

[109] P. Billoir, “Progressive track recognition with a Kalman like fitting procedure”, *Comput. Phys. Commun.* **57** (1989) 390–394, doi:10.1016/0010-4655(89)90249-X.

[110] P. Billoir and S. Qian, “Simultaneous pattern recognition and track fitting by the Kalman filtering method”, *Nucl. Instrum. Meth. A* **294** (1990) 219–228, doi:10.1016/0168-9002(90)91835-Y.

[111] R. Fruhwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Meth. A* **262** (1987) 444–450, doi:10.1016/0168-9002(87)90887-4.

[112] S. Lantz et al., “Speeding up particle track reconstruction using a parallel Kalman filter algorithm”, *JINST* **15** (2020), no. 09, P09030, doi:10.1088/1748-0221/15/09/P09030, arXiv:2006.00071.

[113] CMS Collaboration, “Performance of Run 3 track reconstruction with the mkFit algorithm”, CMS Detector Performance Summary, 2022. CMS-DP-2022-018.

[114] CMS Collaboration, “Performance of the track selection DNN in Run 3”, CMS Detector Performance Summary, 2023. CMS-DP-2023-009.

[115] CMS Collaboration, “2017 tracking performance plots”, CMS Detector Performance Summary, 2017. CMS-DP-2017-015.

- [116] CMS Collaboration, “Muon tracking performance in the CMS Run-2 Legacy data using the tag-and-probe technique”, CMS Detector Performance Summary, 2020. CMS-DP-2020-035.
- [117] CMS Collaboration, “Tracking performances for charged pions with Run2 Legacy data”, CMS Detector Performance Summary, 2022. CMS-DP-2022-012.
- [118] W. Waltenberger, R. Frühwirth, and P. Vanlaer, “Adaptive vertex fitting”, *Journal of Physics G: Nuclear and Particle Physics* **34** (nov, 2007) N343, doi:10.1088/0954-3899/34/12/N01.
- [119] CMS Collaboration, “Displaced tracking and vertexing calibration using neutral K mesons”, CMS Detector Performance Summary, 2024. CMS-DP-2024-010.
- [120] W. Adam, R. Frühwirth, A. Strandlie, and T. Todorov, “Reconstruction of electrons with the Gaussian-sum filter in the CMS tracker at the LHC”, *Journal of Physics G: Nuclear and Particle Physics* **31** (July, 2005) N9–N20, doi:10.1088/0954-3899/31/9/n01.
- [121] CMS Collaboration, “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”, *JINST* **16** (2021), no. 05, P05014, doi:10.1088/1748-0221/16/05/P05014, arXiv:2012.06888.
- [122] CMS Collaboration, “Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **05** (2023) 228, doi:10.1007/JHEP05(2023)228, arXiv:2205.08582.
- [123] CMS Collaboration, “Search for long-lived particles decaying to final states with a pair of muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, *JHEP* **05** (2024) 047, doi:10.1007/JHEP05(2024)047, arXiv:2402.14491.

[124] CMS Collaboration, “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018), no. 06, P06015,
doi:10.1088/1748-0221/13/06/P06015,
arXiv:1804.04528.

[125] CMS Collaboration, “Muon reconstruction performance during Run II”, CMS Detector Performance Summary, 2019. CMS-DP-2019-022.

[126] CMS Collaboration, “Muon Reconstruction and Identification Performance with Run-2 data”, CMS Detector Performance Summary, 2020. CMS-DP-2020-040.

[127] CMS Collaboration, “Search for Neutral Higgs Bosons Decaying to Tau Pairs in pp Collisions at $\sqrt{s} = 7$ TeV”, *Phys. Lett. B* **713** (2012) 68–90,
doi:10.1016/j.physletb.2012.05.028,
arXiv:1202.4083.

[128] G. P. Salam, “Towards Jetography”, *Eur. Phys. J. C* **67** (2010) 637–686,
doi:10.1140/epjc/s10052-010-1314-6,
arXiv:0906.1833.

[129] CMS Collaboration, “Jet algorithms performance in 13 TeV data”, CMS Preliminary Analysis Summary, 2017. CMS-PAS-JME-16-003.

[130] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_t jet clustering algorithm”, *JHEP* **04** (2008) 063,
doi:10.1088/1126-6708/2008/04/063,
arXiv:0802.1189.

[131] D. Bertolini, P. Harris, M. Low, and N. Tran, “Pileup Per Particle Identification”, *JHEP* **10** (2014) 059,
doi:10.1007/JHEP10(2014)059, arXiv:1407.6013.

[132] CMS Collaboration, “Pileup mitigation at CMS in 13 TeV data”, *JINST* **15** (2020), no. 09, P09018,

doi:10.1088/1748-0221/15/09/P09018,
arXiv:2003.00503.

[133] CMS Collaboration, “Jet energy scale and resolution measurement with Run 2 Legacy Data Collected by CMS at 13 TeV”, CMS Detector Performance Summary, 2021. CMS-DP-2021-033.

[134] CMS Collaboration, “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”, *JINST* **12** (2017), no. 02, P02014, doi:10.1088/1748-0221/12/02/P02014,
arXiv:1607.03663.

[135] CMS Collaboration, “Performance summary of AK4 jet b tagging with data from proton-proton collisions at 13 TeV with the CMS detector”, CMS Detector Performance Summary, 2023. CMS-DP-2023-005.

[136] CMS Collaboration, “Performance summary of AK4 jet charm tagging with the CMS Run2 Legacy dataset”, CMS Detector Performance Summary, 2023. CMS-DP-2023-006.

[137] CMS Collaboration, “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector”, *JINST* **14** (2019), no. 07, P07004,
doi:10.1088/1748-0221/14/07/P07004,
arXiv:1903.06078.

[138] CMS Collaboration, “Jet Energy Scale and Resolution Measurements Using Run3 Data Collected by the CMS experiment in 2024 at 13.6 TeV”, CMS Detector Performance Summary, 2025.

[139] CMS Collaboration, “Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS”, *JINST* **6** (2011) P11002,
doi:10.1088/1748-0221/6/11/P11002,
arXiv:1107.4277.

[140] CMS Collaboration, “CMS PYTHIA 8 colour reconnection tunes based on underlying-event data”, arXiv:2205.02905.

- [141] CMS Collaboration, “Performance of the pile up jet identification in CMS for Run 2”, CMS Detector Performance Summary, 2020. CMS-DP-2020-020.
- [142] CMS Collaboration, “CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV”, CMS Physics Analysis Summary, 2019. CMS-PAS-LUM-18-002.
- [143] CMS Collaboration, “Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS”, *Eur. Phys. J. C* **81** (2021), no. 9, 800, doi:10.1140/epjc/s10052-021-09538-2, arXiv:2104.01927.
- [144] X. Cid Vidal, Y. Tsai, and J. Zurita, “Exclusive displaced hadronic signatures in the LHC forward region”, *JHEP* **01** (2020) 115, doi:10.1007/JHEP01(2020)115, arXiv:1910.05225.
- [145] E. Bagnaschi, G. Degrassi, P. Slavich, and A. Vicini, “Higgs production via gluon fusion in the POWHEG approach in the SM and in the MSSM”, *JHEP* **02** (2012) 088, doi:10.1007/JHEP02(2012)088, arXiv:1111.2854.
- [146] S. Alioli, P. Nason, C. Oleari, and E. Re, “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”, *JHEP* **06** (2010) 043, doi:10.1007/JHEP06(2010)043, arXiv:1002.2581.
- [147] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with Parton Shower simulations: the POWHEG method”, *JHEP* **11** (2007) 070, doi:10.1088/1126-6708/2007/11/070, arXiv:0709.2092.
- [148] P. Nason, “A New method for combining NLO QCD with shower Monte Carlo algorithms”, *JHEP* **11** (2004) 040, doi:10.1088/1126-6708/2004/11/040, arXiv:hep-ph/0409146.

- [149] CMS Collaboration, “Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements”, *Eur. Phys. J. C* **80** (2020) 4, doi:10.1140/epjc/s10052-019-7499-4, arXiv:1903.12179.
- [150] K. Hamilton, P. Nason, E. Re, and G. Zanderighi, “NNLOPS simulation of Higgs boson production”, *Journal of High Energy Physics* **2013** (Oct, 2013) 222, doi:10.1007/JHEP10(2013)222.
- [151] CMS Collaboration, “Data Scouting and Data Parking with the CMS High level Trigger”, *PoS EPS-HEP2019* (2020) 139, doi:10.22323/1.364.0139.
- [152] K. S. Kolbig and B. Schorr, “A Program Package for the Landau Distribution”, *Comput. Phys. Commun.* **31** (1984) 97–111, doi:10.1016/0010-4655(84)90085-7. [Erratum: *Comput.Phys.Commun.* 178, 972 (2008)].
- [153] CMS Collaboration, “Search for light scalar particles from Higgs boson decays in exclusive final states with two muons and two hadrons”, cms preliminary analysis summary, 2025. CMS-PAS-EXO-24-034.
- [154] CMS Collaboration, “The CMS statistical analysis and combination tool: COMBINE”, *Comput. Softw. Big Sci.* **8** (2024) 19, doi:10.1007/s41781-024-00121-4, arXiv:2404.06614.
- [155] CMS Collaboration, “CMS Luminosity Measurements for the 2016 Data Taking Period”, CMS Physics Analysis Summary, 2017. CMS-PAS-LUM-17-001.
- [156] CMS Collaboration, “CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13$ TeV”, CMS Physics Analysis Summary, 2018. CMS-PAS-LUM-17-004.
- [157] Particle Data Group Collaboration Collaboration, “Review of particle physics”, *Phys. Rev. D* **110** (Aug, 2024) 030001, doi:10.1103/PhysRevD.110.030001.

- [158] A. L. Read, “Presentation of search results: the CLs technique”, *Journal of Physics G: Nuclear and Particle Physics* **28** (sep, 2002) 2693, doi:10.1088/0954-3899/28/10/313.
- [159] CMS Collaboration, “Long-lived particle triggers at CMS: Strategy and performance in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, CMS Physics Analysis Summary, 2025. CMS-PAS-EXO-23-016.
- [160] CMS Collaboration, “The Phase-2 Upgrade of the CMS Tracker”, technical design report, 2017. doi:10.17181/CERN.QZ28.FLHW, CERN-LHCC-2017-009, CMS-TDR-014.
- [161] CODEX-b Collaboration, “CODEX-b: Opening New Windows to the Long-Lived Particle Frontier at the LHC: ESPP Update Comprehensive summary”, arXiv:2505.05952.
- [162] MATHUSLA Collaboration, “MATHUSLA: An External Long-Lived Particle Detector to Maximize the Discovery Potential of the HL-LHC”, arXiv:2504.01999.
- [163] M. Bauer, O. Brandt, L. Lee, and C. Ohm, “ANUBIS: Proposal to search for long-lived neutral particles in CERN service shafts”, arXiv:1909.13022.
- [164] PBC Collaboration, “Summary Report of the Physics Beyond Colliders Study at CERN”, arXiv:2505.00947.