

Low Energy Leptons in High Energy Physics

Low Energy Leptons in High Energy Physics

Peter Meiring

A Search for Physics beyond the Standard Model with
Compressed Mass-Spectra and New Algorithms for
Triggering on Electrons at the High Luminosity LHC

Peter Meiring

CERN-THESIS-2024-115
27/03/2024

Low Energy Leptons in High Energy Physics

A Search for Physics beyond the Standard Model with Compressed Mass-Spectra and New Algorithms for Triggering on Electrons at the High Luminosity LHC

Dissertation

zur

Erlangung der naturwissenschaftlichen Doktorwürde
(Dr. sc. nat.)

vorgelegt der

Mathematisch-naturwissenschaftlichen Fakultät

der

Universität Zürich

von

Peter Eduard Meiring

aus

den Niederlanden

Promotionskommission

Prof. Dr. Cristina Botta (Vorsitz)

Prof. Dr. Florencia Canelli

Dr. Maurizio Pierini

Prof. Dr. Nicola Serra

Zürich, 2024

**Low Energy Leptons in High Energy Physics – A Search for Physics
beyond the Standard Model with Compressed Mass-Spectra and New Algorithms
for Triggering on Electrons at the High Luminosity LHC**



**University of
Zurich**^{UZH}

The research presented in this thesis was performed at the Physics Institute of the University of Zurich (UZH), Switzerland. It was supervised by Prof. Dr. Cristina Botta and Dr. Jaana Heikkilä.

Cover design by Cynthia Meiring

Printed in the Netherlands by Drukwerkknodig.nl

It's a matter of definition ...

... and perspective

Acknowledgements

Thank you to Cristina and Jaana, for your exceptional supervision and mentorship. And to Paris, Marco, Christian, Gianluca and Sioni.

Thank you to my friends at CERN; Sergio, Marina, Umberto, Carlos, Santeri, Celia, Izaak, Kyle, Tiziano. And to the CMS group at UZH.

Thank you to my flatmates in Saint-Genis; Giuseppe, Rahul, Daniele, Manos, Miguel, Ricardo.

Thank you to my close ones; My family and Huiling.

Publications

- [I] CMS Collaboration, “Search for supersymmetry in final states with two or three soft leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2022) 091, [doi:10.1007/JHEP04\(2022\)091](https://doi.org/10.1007/JHEP04(2022)091), [arXiv:2111.06296](https://arxiv.org/abs/2111.06296)
- [II] CMS Collaboration, “Combined search for electroweak production of winos, binos, higgsinos, and sleptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **109** (2024) 112001, [doi:10.1103/PhysRevD.109.112001](https://doi.org/10.1103/PhysRevD.109.112001), [arXiv:2402.01888](https://arxiv.org/abs/2402.01888).
- [III] P.E. Meiring for the CMS Collaboration, “Constraining challenging regions of the SUSY parameter space with the CMS experiment”, *PoS EPS-HEP2021* (2022) 659, [doi:10.22323/1.398.0659](https://doi.org/10.22323/1.398.0659)
- [IV] P.E. Meiring for the CMS Collaboration, “Triggering on electrons, photons, tau leptons, jets and energy sums at HL-LHC with the upgraded CMS Level-1 Trigger”, *PoS ICHEP2022* (2023) 939, [doi:10.22323/1.414.0939](https://doi.org/10.22323/1.414.0939)
- [V] CMS Collaboration, “Electron reconstruction and identification in the CMS Phase-2 Level-1 Trigger”, *CMS Detector Performance Note CMS-DP-2023-047*, 2023

Author’s Contributions

The main results of this thesis are reported in Chapters 5, 6 and 7.

For the search reported in Chapter 5, the author played a leading role as main analyzer and contact person on behalf of the analysis team within the CMS Collaboration. The author also contributed to the preceding search, reported in Section 5.2, by improving the signal modeling, optimizing event selections to extend the analysis acceptance to lower signal mass-splittings, and developing the parametric signal extraction procedure.

For the combination effort of Chapter 6, the author had a central position in the analysis team, organizing the analysis inputs and infrastructure for the combination, performing studies of uncertainty correlation schemes and executing and validating the full combination itself. The author also served as contact person on behalf of the “ $2/3\ell$ soft” analysis (ie. that of Section 5.2) within the combination team and added the new slepton interpretation in the mass-compressed region of the model parameter space.

For the Phase-2 Level-1 trigger work reported in Chapter 7, the author was the main developer of the *composite electron identification* strategy, performing simulation-based feasibility studies, training the identification model with machine learning techniques and contributing to the implementation of the trigger algorithm in the emulator and firmware of the Level-1 Trigger upgrade.

Finally, the author also contributed to the data-taking effort of the CMS experiment during LHC Run 2, by serving as online *and* offline shifter for the Level-1 Trigger system. This includes monitoring and data certification tasks.

Contents

Introduction	5
1 The Standard Model of Particle Physics	5
1.1 Overview	6
1.2 The Standard Model Lagrangian	8
1.2.1 Gauge Term	8
1.2.2 Fermion Term	9
1.2.3 Higgs Term	11
1.2.4 Yukawa Term	12
1.3 Spontaneous Symmetry Breaking	12
1.4 Quantum Chromodynamics	14
1.4.1 Running Coupling and Renormalization	14
1.4.2 Factorization and Parton Distribution Functions	17
1.5 Beyond the Standard Model	17
2 Supersymmetry	21
2.1 Motivation for Supersymmetry at the TeV-Scale	22
2.2 The Minimal Supersymmetric Standard Model	25
2.2.1 Theoretical Structure	25
2.2.2 Particle Content	29
2.3 Searches for Supersymmetry at the LHC	30
2.3.1 Simplified Models	31
2.3.2 Experimental Search Strategies	32
2.3.3 Run 2 Results from the CMS Experiment	34
2.4 Supersymmetry with Unconventional Signatures	37
2.4.1 Compressed Mass-Spectra	38
2.4.2 Long-lived Particles	39

Contents

3 The LHC and the CMS Experiment	43
3.1 The Large Hadron Collider	44
3.2 The CMS Detector	46
3.2.1 The Solenoid Magnet	49
3.2.2 The Inner Tracker	50
3.2.3 The Electromagnetic Calorimeter	52
3.2.4 The Hadronic Calorimeter	53
3.2.5 The Muon system	54
3.3 The CMS Trigger System	56
3.3.1 The Level-1 Trigger	57
3.3.2 The High Level Trigger	63
3.4 Beyond the LHC: Phase 2	64
3.4.1 The High Luminosity LHC	65
3.4.2 The Phase 2 Upgrade of the CMS Detector	67
4 Object and Event Reconstruction in CMS	73
4.1 Particle Flow	74
4.1.1 KF Tracks and Vertices	75
4.1.2 Calorimeter Clusters	80
4.1.3 PF Element Linking Scheme	81
4.2 Muons	82
4.3 Electrons and Isolated Photons	84
4.3.1 Mustache Superclusters	86
4.3.2 GSF Tracking	86
4.3.3 Refined Superclusters and Conversion Finding	88
4.3.4 GED and PF e/γ Objects	89
4.3.5 LowPt Electrons	91
4.4 Hadrons and Non-isolated Photons	93
4.5 Higher-level Objects	93
4.5.1 Jets and Flavour Tagging	94
4.5.2 Hadronic Taus	96
4.5.3 Energy Sums	99
4.6 Pileup Per Particle Identification	101
5 Search for Mass-compressed Prompt and Long-lived Electroweakinos with Soft Lepton Final States	105
5.1 Signal Models	106
5.2 Preceding Research in CMS	108

5.2.1	Legacy Run 2 Analysis	108
5.2.2	Parametric Signal Extraction	112
5.3	Overview of the New Analysis Strategy	116
5.4	Data and Simulated Samples	118
5.4.1	Trigger Strategy	118
5.4.2	Monte-Carlo Simulations	123
5.5	Object Definitions	129
5.5.1	Prompt Electrons	130
5.5.2	Prompt Muons	138
5.5.3	Displaced Dimuons	141
5.5.4	Jets and Missing Transverse Momentum	147
5.6	Event Selection and Categorization	149
5.6.1	Prompt Search Regions	151
5.6.2	Displaced Search Regions	152
5.7	Background Estimation	153
5.7.1	Overview	153
5.7.2	Prompt Lepton Processes	154
5.7.3	Non-prompt or Fake Leptons	160
5.7.4	Displaced Lepton Processes	167
5.8	Systematic Uncertainties	171
5.9	Results	173
5.9.1	Binning optimization	178
5.9.2	Signal Interpretations	179
5.10	Conclusion	183
6	Combination of Searches for Electroweak Supersymmetry	187
6.1	Pushing the Limits with Statistical Combination	188
6.2	Signal Models	189
6.3	Individual Searches	192
6.3.1	(Semi)leptonic Final States	192
6.3.2	All-hadronic Final States	196
6.4	Combination Strategy	198
6.4.1	Orthogonality of Search Regions	199
6.4.2	Systematic Uncertainties	200
6.5	Results and Interpretations	202
6.6	Conclusion	204

Contents

7 Level-1 Trigger Algorithms for Electrons at the HL-LHC	211
7.1 The Phase 2 Upgrade of the CMS Level-1 Trigger	212
7.2 Track-matched Electron Reconstruction in the CMS Endcaps	216
7.2.1 HGCAL Trigger Primitive Generation	216
7.2.2 Stand-alone e/ γ Identification	218
7.2.3 Track Trigger Primitive Generation	219
7.2.4 Track-matched Electron Identification	222
7.3 Improving Performance for Track-matched Electrons	224
7.3.1 Composite Electron Identification	224
7.3.2 Model Training	228
7.3.3 Performance in simulation	232
7.3.4 Firmware implementation	237
7.4 Conclusion	245
Summary	253
A Limit Setting in Searches for New Physics at the LHC	253
B Data-driven Estimation of Non-prompt or Fake Leptons	257
B.1 Fake Rate Measurements	257
B.2 Transfer Factor Derivations	265
B.3 Semi-data-driven Scale Factors	267
C Features for Multi-variate Electron Identification in the CMS Phase-2 Level-1 Trigger	273
Bibliography	300

Introduction

The CERN Large Hadron Collider (LHC) is the world's largest and most powerful particle accelerator. It consists of a 27 km long superconducting circular collider that is located 100 meters underground, just west of Lake Geneva. The LHC is designed to probe the frontiers of energy and luminosity, hence serves as the primary tool to study properties of the Standard Model (SM) [1–5] of particle physics and search for new physics phenomena. The four main detectors along its circumference are ATLAS, CMS, ALICE and LHCb, which together pursue a broad and diverse physics program. Since the start of operations in 2010, this has lead to numerous achievements, among which most notably the discovery of the Higgs boson by ATLAS and CMS in 2012 [6, 7] - a scalar particle predicted by the SM theory.

With the data collected by the end of LHC Run 2 in 2018, the experiments have validated SM predictions over more than 10 orders of magnitude. However, there are still several fundamental questions related to eg. gravity, the hierarchy problem and the origin of Dark Matter, to which the SM does not formulate an answer. Among the many proposed beyond-the-SM (BSM) theories, Supersymmetry [8–12] is arguably the most remarkable one, suggesting the existence of new particle states at the electroweak scale (100 GeV - 1 TeV), ie. within the LHC discovery potential. However, no experimental evidence has been found so far.

In the absence of hints of new physics phenomena at the electroweak scale, the general tendency is to assume that these phenomena instead reside at higher energy scales, which cannot be probed with current collider technologies. It should be noted however, that new physics could *still* reside at the electroweak scale through realizations that are experimentally harder to probe. Search strategies therefore continuously develop in broadly three directions:

- **New searches with unconventional signatures** – Experimental observations have excluded electroweak scale BSM physics realizations in their most trivial form. In collision events these would yield final state signa-

Introduction

tures that are relatively easy to distinguish from SM processes. However, new physics phenomena *could* have been missed in previous searches if they display themselves through unconventional signatures that are more challenging to capture experimentally. Examples include new particle states with *compressed mass-spectra*, or realizations that result in soft unclustered energy patterns, both of which are hard to distinguish in the overwhelming hadronic background of the LHC environment. Furthermore, several BSM theories predict the presence of new particle states that are *long-lived* and decay at a notable distance from the primary proton-proton interaction point. Since standard particle reconstruction techniques are instead designed for objects with provenance from the primary vertex, these long-lived particles *could* have escaped experimental observation. As such, these examples correspond to corners of the parameter space that are hardly probed, but *must* be explored to give the ultimate answer to the presence of new physics at the electroweak scale.

- **Maximizing analysis sensitivity** – Analysis strategies are continuously improved, eg. with the usage of machine learning techniques, to maximally exploit the amount of data collected so far. Maximizing the sensitivity is crucial, not only for searches of BSM physics, but also for precision measurements of SM properties, which could serve as probe for new physics as well through potential deviations from the SM expectation. Besides improving the strategies of individual analysis efforts, the sensitivity can furthermore be enhanced via statistical combinations of multiple analyses that target different production or decay channels.
- **Collecting more data** – Precision measurements are typically limited by *systematic* uncertainties, that can be reduced via a better understanding of theory and experiment. On the other hand, BSM physics searches are generally limited by *statistical* uncertainties, that can only be decreased by collecting more collision data. As such, the LHC continues to deliver collision data during Run 3 (2022-2025) and will furthermore be upgraded to the High Luminosity LHC (HL-LHC) from 2026 onwards, which will operate at a much higher instantaneous luminosity and thereby significantly shrink the statistical uncertainties.

The work presented in this thesis contributes to the exploration of the *high energy* electroweak scale in each of these directions, with as a common thread the usage of *low energy* leptons. The thesis is structured as follows.

Chapter 1 gives a theoretical overview of the SM theory and its fundamental shortcomings. As mentioned above, Supersymmetry is a particularly attractive extension that addresses many of these shortcomings and predicts the presence a supersymmetric partner for each SM particle. Besides a brief description of the framework of Supersymmetry, Chapter 2 mainly focuses on motivating the existence of superpartners at the electroweak scale and the experimental search strategies at the LHC.

Chapters 3 and 4 provide a description of the LHC itself and the CMS experiment, as well as the reconstruction techniques for physics objects in collision events at CMS. An emphasis is put on the reconstruction of electrons, particularly those with low momentum, as these play a central role in the work described in the remainder of the thesis.

Chapter 5 presents a search for new physics phenomena with unconventional signatures. In particular, the search targets realizations of BSM physics with compressed mass-spectra using final states with multiple low energy (soft) leptons. A new reconstruction technique is employed, dedicated to soft electrons, to probe mass-splittings down to the sub-GeV level. The analysis furthermore targets scenarios in which new particle states are mass-compressed *and* long-lived as motivated by models of Supersymmetry. These topologies lead to final states of soft *and* displaced leptons, that have not been probed before.

The search of Chapter 5 is used as one of the input analyses for the legacy Run 2 combination of searches for electroweak Supersymmetry at CMS, which is described in Chapter 6. Searches targeting both (semi-)leptonic and hadronic final states are combined to maximize the sensitivity to potential supersymmetric realizations. New interpretations are made for scenarios with mass-compressed sleptons (the superpartners of SM leptons), making use of novel signal extraction techniques.

Finally, Chapter 7 presents work in the context of the real-time CMS data-taking at the future HL-LHC. The expected increase of instantaneous luminosity of the HL-LHC poses serious challenges for the experiment, hence considerable upgrades are needed to continue operations and maintain the physics acceptance of Run 2 also in the HL-LHC data-taking conditions. As such, the system responsible for the real-time event selection (the Level-1 Trigger) has been completely redesigned. In this chapter, a new algorithm is presented for the real-time identification of (soft) electrons, making use of machine-learning techniques and exploiting the more granular information from sub-detector upgrades and new trigger capabilities.

Chapter 1

The Standard Model of Particle Physics

The Standard Model (SM) of particle physics is a theory that describes elementary particles and their interactions. It has been developed during the past century through a series of experimental observations and concepts proposed by outstanding theorists, and has established itself as one of the greatest achievements in physics. Describing the full SM theory goes beyond the scope of this thesis. Instead, this chapter provides a condensed summary, exposing the free parameters of the theory that have to be measured experimentally, and highlighting some of the relevant concepts for the physics processes studied at the Large Hadron Collider. Shortcomings of the SM, which motivate the searches for new physics beyond the SM described in Chapters 5 and 6, are discussed as well. The material is largely based on Refs. [13–15].

The Standard Model of Particle Physics

1.1 Overview

The Standard Model (SM) of particle physics [1–5] is a theoretical framework that describes the fundamental building blocks of matter and their interactions. It is formulated in the language of group theory and quantum field theory (QFT), and is known for its remarkable predictive power.

The SM particle content, shown in Figure 1.1, can be divided into two categories based on the spin quantum numbers. Particles with half-integer spin are referred to as *fermions*, which can be further split into *leptons* and *quarks*. Leptons are susceptible only to the electromagnetic and weak forces, while quarks are subject also to the strong force. The forces are mediated by particles with integer spin, referred to as *gauge bosons*. The W- and Z-bosons carry the weak force, the photon (γ) carries the electromagnetic force and the gluons are carriers of the strong force. The Higgs boson is the only particle with spin 0 and plays a central role in spontaneous symmetry breaking, which generates the particle masses. Gravity is not accounted for in the SM, and is relevant only for vastly larger mass and distance scales.

The underlying symmetry structure that dictates the way that particles interact, is given by the gauge group $G = SU(3)_c \times SU(2)_L \times U(1)_Y$, which can be split



Figure 1.1: Particle content of the Standard Model. Within each box, the values indicate from top to bottom: particle mass, charge and spin [16].

1.1 Overview

in two parts. Firstly, $SU(2)_L \times U(1)_Y$ is a chiral gauge group¹ that corresponds to the electroweak part of the SM. The $SU(2)_L$ group has three gauge bosons ($W^i, i \in \{1,2,3\}$, one for each of its generators), gauge coupling g , and acts on the flavor indices of the left-handed fermion components. The $U(1)_Y$ group has one gauge boson (B) and coupling g' . The weak and electromagnetic interactions were unified by Glashow, Salam and Weinberg as $SU(2)_L \times U(1)_Y$ that is spontaneously broken into $U(1)_Q$. Through this spontaneous symmetry breaking (SSB) mechanism, the W^i and B mix into the massive W^\pm and Z bosons, as well as the massless photon, hence incorporating quantum electrodynamics (QED) in the theory of weak interactions. Secondly, $SU(3)_c$ is a non-chiral gauge group that corresponds to the strong part of the SM. It has eight bosons ($G^i, i \in \{1\dots8\}$), gauge coupling g_s and acts on the color indices of the quarks. Unlike $SU(2)_L \times U(1)_Y$, the $SU(3)_c$ group is not spontaneously broken, implying that the associated bosons (gluons) remain massless.

A key feature of the SM theory is the fact that it is a *renormalizable* QFT. This means that non-physical infinities arising from calculations of particle interactions can be absorbed by redefinition of a finite number of parameters of the theory. An example of this will be shown in Section 1.4.1. There are 19 free parameters that have to be experimentally measured in order to fully define the theory and form predictions of the interactions at arbitrary energy scales:

- 3 gauge couplings: g_s (or alternatively $\alpha_s = g_s^2/4\pi$), g and g' . The latter two can be parameterized in terms of the Weinberg angle θ_W (which describes the mixing of the W^3 and B bosons into the γ and Z boson) and fine structure constant α , namely: $\tan(\theta_W) = g'/g$ and $\alpha = e^2/4\pi$, where $e = g \sin(\theta_W)$, using natural units²
- 6 quark Yukawa couplings (or masses)
- 3 charged lepton Yukawa couplings (or masses)
- 3 angles and 1 complex phase of the Cabibbo-Kobayashi-Maskawa (CKM) matrix, corresponding to the mixing of quark flavor eigenstates into mass eigenstates and the charge-parity (CP) violation in the weak decays of quarks
- 1 QCD vacuum angle, θ_{QCD} , corresponding to CP violation in the strong interaction

¹A chiral gauge group treats left- and right-handed components of particles differently.

²The convention $\epsilon_0 = c = \hbar = 1$ is also used throughout this thesis.

The Standard Model of Particle Physics

- 2 Higgs parameters: the mass m_h and vacuum expectation value v . The latter is related to the tree-level masses of the W^\pm and Z bosons as $M_W = vg/2$ and $M_Z = v\sqrt{g^2 + g'^2}/2$

A minimal extension of the SM incorporates massive neutrinos, facilitating the observed neutrino oscillations [17] by adding 7 more free parameters:

- 3 neutral lepton Yukawa couplings (or masses)
- 3 angles and 1 complex phase of the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, corresponding to the mixing of neutrino flavor and mass eigenstates, and CP violation in the neutrino sector

1.2 The Standard Model Lagrangian

The particle dynamics are described with the *SM Lagrangian density* (or simply *SM Lagrangian*), which adheres to the symmetry structure of the $SU(3)_c \times SU(2)_L \times U(1)_Y$ gauge group. It can be factorized³ as:

$$\mathcal{L} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{fermion}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}} \quad (1.1)$$

where the subscripts refer to the various sectors of the theory, as described below.

1.2.1 Gauge Term

The gauge term of the SM Lagrangian is expanded as:

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}G_{\mu\nu}^i G^{\mu\nu i} - \frac{1}{4}W_{\mu\nu}^i W^{\mu\nu i} - \frac{1}{4}B_{\mu\nu} B^{\mu\nu} \quad (1.2)$$

Where $G_{\mu\nu}^i$, $W_{\mu\nu}^i$ and $B_{\mu\nu}$ are the field strength tensors associated to the gluons ($SU(3)_c$) and W and B ($SU(2)_L \times U(1)_Y$) fields (groups), respectively, defined as:

$$G_{\mu\nu}^i = \partial_\mu G_\nu^i - \partial_\nu G_\mu^i - g_s f_{ijk} G_\mu^j G_\nu^k \quad i, j, k \in \{1 \dots 8\} \quad (1.3)$$

$$W_{\mu\nu}^i = \partial_\mu W_\nu^i - \partial_\nu W_\mu^i - g \epsilon_{ijk} W_\mu^j W_\nu^k \quad i, j, k \in \{1, 2, 3\} \quad (1.4)$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \quad (1.5)$$

³Ghost and gauge fixing terms are omitted. These are normally introduced to cancel artifacts that may arise when quantizing gauge theories, hereby maintaining unitarity (ie. conservation of probability) and renormalizability of the theory. Gauge fixing terms constrain a gauge symmetry to eliminate certain associated degrees of freedom.

1.2 The Standard Model Lagrangian

Besides the kinetic terms in each of the above equations, the last terms in Equations 1.3 and 1.4 (with structure constants f_{ijk} and ϵ_{ijk}) reveal the *non-abelian* (ie. non-commutative) structure of the $SU(3)_c$ and $SU(2)_L$ groups, and consequently give rise to the three- and four-point self-interactions of G^i and W^i . No mass-terms are visible yet, as these are obtained only after SSB when the W^i and B fields mix into the W^\pm , Z and γ bosons. Finally, a term $\theta_{QCD} g_s^2 G_{\mu\nu}^i \tilde{G}^{\mu\nu i} / 32\pi^2$ with $\tilde{G}_{\mu\nu}^i$ the dual field strength tensor, may be added to the Lagrangian, corresponding to CP violation in the strong interaction for $\theta_{QCD} \neq 0$.

1.2.2 Fermion Term

The fermion sector consists of three families, each with five fields, being a combination of left-handed $SU(2)$ doublets and right-handed $SU(2)$ singlets:

$$q_{mL}^0 = \begin{pmatrix} u_m^0 \\ d_m^0 \end{pmatrix}_L, \ell_{mL}^0 = \begin{pmatrix} \nu_m^0 \\ e_m^{-0} \end{pmatrix}_L, u_{mR}^0, d_{mR}^0, e_{mR}^{-0} \quad (1.6)$$

The subscripts refer to the family $m \in \{1, 2, 3\}$ and chirality, while the superscript refers to the fact that these are weak eigenstates. After SSB the u^0 and d^0 (ν^0 and e^{-0}) are mixed into the quark (lepton) mass eigenstates presented in Figure 1.1. Adding another right-handed $SU(2)$ singlet ν_{mR}^0 allows for massive neutrinos in many SM extensions.

The way that the fermions interact with the gauge bosons is characterized with the quantum numbers associated to the $SU(3)_c$, $SU(2)_L$ and $U(1)_Y$ groups. The generators of $SU(2)_L$ are the components of the *weak isospin* operator $T^i = \tau^i/2$, with $i \in \{1, 2, 3\}$ and τ^i representing the (2×2) Pauli matrices. The eigenvalue of the third component (T^3) acting on a left-handed $SU(2)_L$ doublet is the weak isospin quantum number t_L^3 . Right-handed fermions are singlets in $SU(2)_L$, and as such $t_L^3 = 0$ and these particles do not participate in weak isospin interactions. Analogously, the generators of $SU(3)_c$ are the components of the *color charge* operator $\Lambda^i = \lambda^i/2$, with $i \in \{1 \dots 8\}$ and λ^i denoting the (3×3) Gell-Mann matrices, which acts only on color triplets. Finally, the generator of $U(1)_Y$ is the *weak hypercharge* operator Y with associated quantum number y . The representations and quantum numbers for the SM fields (plus $\nu_{m,R}^0$) is shown in Table 1.1.

The Standard Model of Particle Physics

Field	Representations		Quantum numbers		
	$SU(3)_c$	$SU(2)_L$	t_L^3	y	$q = t_L^3 + y$
$q_{mL} = \begin{pmatrix} u_m \\ d_m \end{pmatrix}_L$	3	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$+\frac{1}{6}$	$\begin{pmatrix} +\frac{2}{3} \\ -\frac{1}{3} \end{pmatrix}$
u_{mR}	3	1	0	$+\frac{2}{3}$	$+\frac{2}{3}$
d_{mR}	3	1	0	$-\frac{1}{3}$	$-\frac{1}{3}$
$\ell_{mL} = \begin{pmatrix} \nu_m \\ e_m^- \end{pmatrix}_L$	1	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$-\frac{1}{2}$	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$
e_{mR}^-	1	1	0	-1	-1
ν_{mR}	1	1	0	0	0
G^i	8	1	0	0	0
$W^i = \begin{pmatrix} W^+ \\ W^- \\ W^3 \end{pmatrix}$	1	3	$\begin{pmatrix} +1 \\ -1 \\ 0 \end{pmatrix}$	0	$\begin{pmatrix} +1 \\ -1 \\ 0 \end{pmatrix}$
B	1	1	0	0	0
$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$	1	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$+\frac{1}{2}$	$\begin{pmatrix} +1 \\ 0 \end{pmatrix}$
$\tilde{\phi} = \begin{pmatrix} \phi^{0\dagger} \\ \phi^- \end{pmatrix}$	1	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$-\frac{1}{2}$	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$

Table 1.1: Representations of the SM fields (plus ν_{mR}^0) under $SU(3)_c$ and $SU(2)_L$, as well as the quantum numbers t_L^3 (weak isospin) and y (weak hypercharge) associated respectively to $SU(2)_L$ and $U(1)_Y$. Also the quantum number $q = t_L^3 + y$ (electric charge) is shown, corresponding to the $U(1)_Q$ gauge group that emerges after SSB. Bold numbers represent the dimension of the representations: **1**, **2**, **3** and **8** denote a singlet, doublet, triplet and octet, respectively. The W^1 and W^2 fields are related to the mass eigenstates W^+ and W^- as $W^\pm = (W^1 \mp iW^2)/\sqrt{2}$. The superscript ⁰ for the fermion weak eigenstates has been omitted to avoid confusion with the neutral component of the Higgs doublet, while subscript m ($m \in \{1, 2, 3\}$) corresponds to the fermion families.

1.2 The Standard Model Lagrangian

The fermion term of the Lagrangian may now be written as:

$$\mathcal{L}_{\text{fermion}} = \sum_{m=1}^3 (i\bar{q}_{mL}^0 \not{D} q_{mL}^0 + i\bar{\ell}_{mL}^0 \not{D} \ell_{mL}^0 + i\bar{u}_{mR}^0 \not{D} u_{mR}^0 + i\bar{e}_{mR}^0 \not{D} e_{mR}^0 + i\bar{\nu}_{mR}^0 \not{D} \nu_{mR}^0) \quad (1.7)$$

where $\not{D} = \gamma^\mu D_\mu$, with γ^μ denoting the Dirac matrices, and the Dirac adjoint of a fermion field ψ is $\bar{\psi} = \psi^\dagger \gamma^0$. The covariant derivative D_μ is expanded as:

$$D_\mu = \partial_\mu + ig_s \Lambda_j G_\mu^j + ig T_j W_\mu^j + ig' Y B_\mu \quad (1.8)$$

Summation over the components j is implied. As can be seen from Equations 1.7 and 1.8, the fermion term of the Lagrangian includes kinetic terms, as well as interactions between fermions and gauge bosons, where it should be noted that several terms (such as $-\bar{e}_{mR}^0 \gamma^\mu g T_j W_\mu^j e_{mR}^0$) evaluate to zero, as explained above. Fermion masses are generated through the Yukawa couplings between fermions and the Higgs field (see Section 1.2.4).

1.2.3 Higgs Term

The Higgs term of the SM Lagrangian is:

$$\mathcal{L}_{\text{Higgs}} = (D^\mu \phi)^\dagger D^\mu \phi - V(\phi) \quad (1.9)$$

The field ϕ is a complex isospin doublet, presented also in Table 1.1. As can be seen from the table and Equation 1.8, the first term in Equation 1.9 includes a kinetic term as well as the three and four-point interaction between the Higgs and gauge fields. The Higgs potential $V(\phi)$, which maintains $SU(2)_L \times U(1)_Y$ gauge invariance and renormalizability, has the form:

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2 \quad (1.10)$$

where μ and λ are free parameters. That said, λ must be positive for the sake of vacuum stability⁴ and is known as the Higgs self-coupling. For $\mu^2 < 0$, SSB occurs, as described in Section 1.3.

⁴In other words, for $\lambda < 0$ the potential would be unbounded from below.

The Standard Model of Particle Physics

1.2.4 Yukawa Term

The Yukawa coupling between the SM fermions (plus ν_{mR}^0) and the Higgs doublet, needed for the generation of fermion masses through SSB, is given by:

$$\mathcal{L}_{\text{Yukawa}} = - \sum_{m,n=1}^3 (\Gamma_{mn}^u \bar{q}_{mL}^0 \tilde{\phi} u_{nR}^0 + \Gamma_{mn}^d \bar{q}_{mL}^0 \tilde{\phi} d_{nR}^0 + \Gamma_{mn}^e \bar{\ell}_{mL}^0 \tilde{\phi} e_{nR}^0 + \Gamma_{mn}^{\nu} \bar{\ell}_{mL}^0 \tilde{\phi} \nu_{nR}^0) + h.c. \quad (1.11)$$

where $\tilde{\phi} = i\tau^2 \phi^\dagger$ is the conjugate form of the Higgs doublet, also shown in Table 1.1. The Γ^u matrices are arbitrary (3×3) matrices containing the Yukawa couplings, which make up the bulk of free parameters in the SM.

1.3 Spontaneous Symmetry Breaking

Although many of the SM particles observed in experiments are massive, no mass-terms are visible yet in the SM Lagrangian. However, the Higgs field and potential were introduced for this very reason. The complex $SU(2)$ doublet presented in Table 1.1 is in fact the minimal Higgs model that generates the gauge boson masses. In a Hermitian basis it is written as:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix} \quad (1.12)$$

Since the Higgs field has 4 degrees of freedom, the minima of $V(\phi)$ for $\mu^2 < 0$ form a 4-dimensional sphere⁵, satisfying:

$$|\phi|^2 = \frac{-\mu^2}{2\lambda} \equiv \frac{\nu^2}{2} \quad (1.13)$$

where ν is known as the *vacuum expectation value*, and one is free to choose the mean of the Higgs field in the vacuum state as:

$$\langle \phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ \nu \end{pmatrix} \quad (1.14)$$

⁵Often a simplified scenario of a complex scalar $\phi = (\phi_1 + i\phi_2)/\sqrt{2}$ is used, ie. with 2 degrees of freedom, in which case the minima can be visualised as a circle on a *Mexican hat* shaped potential.

1.3 Spontaneous Symmetry Breaking

At this point, it should be noted that the choice of $\mu^2 < 0$ spontaneously breaks each of the generators associated to $SU(2)_L$ and $U(1)_Y$, since $T^i \langle \phi \rangle \neq 0$ and $Y \langle \phi \rangle \neq 0$. However, a new $U(1)_Q$ symmetry emerges, with generator $Q = T^3 + Y$, which leaves the vacuum invariant, since $Q \langle \phi \rangle = 0$. As such, $SU(2)_L \times U(1)_Y \rightarrow U(1)_Q$ will lead to three massive gauge bosons, while the fourth boson is massless.

The physical particle content will become visible by considering small perturbations of the Higgs field around the vacuum state:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \xi^1 + i\xi^2 \\ \nu + h + i\xi^3 \end{pmatrix} \quad (1.15)$$

where ξ^i , $i \in \{1, 2, 3\}$, are Goldstone bosons and h is the physical Higgs boson. Quantization with an appropriate local gauge transformation, known as the *unitary gauge*, causes these non-physical Goldstone bosons to disappear⁶, such that:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ \nu + h \end{pmatrix} \quad (1.16)$$

Expanding the Higgs term of the SM Lagrangian (Equation 1.9) with the Higgs field in unitary gauge (Equation 1.16) will expose the physical gauge bosons and their masses, as well as three- and four-point interactions of the Higgs boson with itself and the other bosons:

$$\begin{aligned} (D^\mu \phi)^\dagger D^\mu \phi &\rightarrow M_W^2 W^\mu+ W^{\mu-} \left(1 + \frac{h}{v}\right)^2 \\ &\quad + \frac{1}{2} M_Z^2 Z^\mu Z_\mu \left(1 + \frac{h}{v}\right)^2 + \frac{1}{2} (\partial_\mu h)^2 \end{aligned} \quad (1.17)$$

$$V(\phi) \rightarrow -\frac{\mu^4}{4\lambda} - \mu^2 h^2 + \lambda v h^3 + \frac{\lambda}{4} h^4 \quad (1.18)$$

where $W_\mu^\pm = (W_\mu^1 \mp iW_\mu^2)/\sqrt{2}$ and $Z_\mu = (-g'B_\mu + gW_\mu^3)/\sqrt{g^2 + g'^2}$. The orthogonal combination of B_μ and W_μ^3 corresponds to the photon⁷. The tree-level masses

⁶This can be seen by using the small angle approximation to write Equation 1.15 as $\phi = e^{i \sum_{i=1}^3 \xi'^i T'^i} (0 \nu + h)^T / \sqrt{2}$, where T'^i are the broken generators T^1 , T^2 and $T^3 - Y$, and using the unitary gauge $\phi \rightarrow \phi' = e^{-i \sum_{i=1}^3 \xi'^i T'^i} \phi$.

⁷The B_μ and W_μ^3 mixing is often parameterized in terms of the Weinberg angle θ_W , with $\tan(\theta_W) = g'/g$, such that $Z_\mu = -\sin(\theta_W)B_\mu + \cos(\theta_W)W_\mu^3$ and the photon field $A_\mu = \cos(\theta_W)B_\mu + \sin(\theta_W)W_\mu^3$.

The Standard Model of Particle Physics

are:

$$M_W = \frac{\nu g}{2}, M_Z = \frac{\nu \sqrt{g^2 + g'^2}}{2}, M_\gamma = 0, M_h = \sqrt{-2\mu^2} \quad (1.19)$$

Finally, SSB also induces the generation of the fermion masses. This can be seen by analogously expanding the Yukawa term (Equation 1.11):

$$\mathcal{L}_{\text{Yukawa}} \rightarrow \bar{u}_L^0 M^u (1 + \frac{h}{\nu}) u_R^0 + (d, e, \nu) \text{ terms} + h.c. \quad (1.20)$$

where the mass matrix $M^u = \Gamma^u \nu / \sqrt{2}$ and $u_L^0 = (u_{1L}^0 \ u_{2L}^0 \ u_{3L}^0)^T$ (and similar for u_R^0). The fermion masses and Yukawa couplings of the fermion mass eigenstates appear when diagonalizing the mass matrices:

$$A_L^{u\dagger} M^u A_R^u = \text{diag}(m_u, m_c, m_t) = \text{diag}(y_u, y_c, y_t) \nu / \sqrt{2} \quad (1.21)$$

$$A_L^{d\dagger} M^d A_R^d = \text{diag}(m_d, m_s, m_b) = \text{diag}(y_d, y_s, y_b) \nu / \sqrt{2} \quad (1.22)$$

$$A_L^{e\dagger} M^e A_R^e = \text{diag}(m_e, m_\mu, m_\tau) = \text{diag}(y_e, y_\mu, y_\tau) \nu / \sqrt{2} \quad (1.23)$$

$$A_L^{\nu\dagger} M^\nu A_R^\nu = \text{diag}(m_{\nu_e}, m_{\nu_\mu}, m_{\nu_\tau}) = \text{diag}(y_{\nu_e}, y_{\nu_\mu}, y_{\nu_\tau}) \nu / \sqrt{2} \quad (1.24)$$

where A_L and A_R are unitary transformations acting on left- and right-handed fermion fields. For the quark sector, these are related to the CKM matrix V_{CKM} , which describes the difference in mixing of the weak and mass eigenstates for the up and down-type quarks, as:

$$V_{CKM} = A_L^{u\dagger} A_R^d \quad (1.25)$$

1.4 Quantum Chromodynamics

As mentioned in Section 1.2.1, the third term in the gluon field strength tensor associated to $SU(3)_c$ (Equation 1.3) gives rise to the three- and four-point self-interactions of the gluons. This has several implications for quantum chromodynamics (QCD) interactions, hence are important in particular for the physics processes studied at the LHC, as will be described in the following.

1.4.1 Running Coupling and Renormalization

Up to now, the couplings corresponding to the gauge groups have been considered constant and without associated energy scale. However, higher-order contributions

1.4 Quantum Chromodynamics

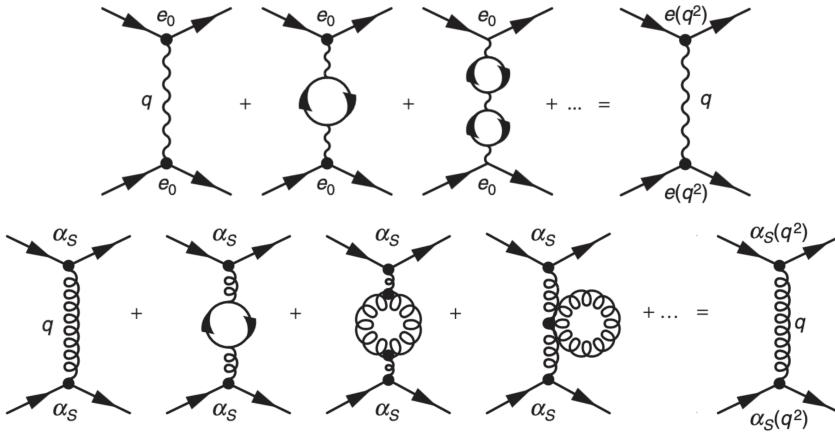


Figure 1.2: Top row: Summation of diagrams contributing to the photon propagator, which carries momentum momentum q , in t-channel fermion interactions. The loops at higher order consist only of fermion pairs. The summation is equivalent to a single photon propagator with an effective running coupling. Bottom row: the QCD analog, where also gluon loops are allowed due to the non-abelian nature of $SU(3)_c$. Figure adapted from [14].

to the boson propagators lead to an experimentally measured coupling strength that depends on the energy scale of the interaction [18].

In QED, when an electrically charged reference particle emits a photon, the higher-order corrections to the photon propagator consist of charged particle pairs, as shown in Figure 1.2 (top row). These vacuum polarization diagrams introduce a Q^2 dependence in the matrix element, where Q is the momentum transfer carried by the photon. An infinite series of such higher-order diagrams exists, leading to non-physical ultraviolet divergences. However, these can be solved by absorbing the Q^2 dependence into an *effective gauge coupling*: $e_0 \rightarrow e(Q^2)$. As a result, the effective QED coupling increases at higher energy scale (or equivalently, lower distance scale).

In contrast to QED, the non-abelian nature of QCD (see the third term in Equation 1.3) gives rise to the three- and four-point gluon self-interactions. Therefore, the higher-order corrections to the gluon propagator consist not only of fermion loops, but also of gluon loops, as shown in Figure 1.2 (bottom row). Due to the color charge carried by the gluons, the gluon loops have the opposite effect with respect to the fermion loops, resulting in an effective QCD coupling that *decreases* with Q^2 (or $1/r^2$). When probed at large enough energy, quarks inside hadrons are therefore behaving more or less as free particles, which is known as *asymptotic freedom*.

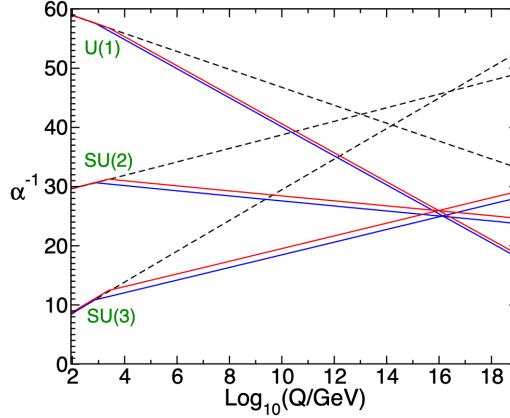


Figure 1.3: Running of the inverse couplings of the gauge groups in the SM (dashed lines) and Minimal Supersymmetric Standard Model (solid lines), corresponding to the two-loop RGE [19].

totic freedom [20, 21]. Consequently, this allows for calculations of quark-qluon interactions with perturbation theory.

For an arbitrary gauge theory the running of the effective coupling g is described by the *renormalization group equation* (RGE):

$$\frac{dg^2}{d\ln Q^2} \equiv 4\pi\beta(g^2) = bg^4 + O(g^6) + \dots \quad (1.26)$$

Using $\alpha = g^2/4\pi$, it is analytically solved at one loop as:

$$\frac{1}{\alpha_s(Q^2)} = \frac{1}{\alpha_2(\mu_R^2)} - 4\pi b \ln\left(\frac{Q^2}{\mu_R^2}\right) \quad (1.27)$$

where μ_R is the chosen *renormalization scale* and the coefficient b is uniquely defined depending on the gauge group. For example, for $SU(3)$ the coefficient $b = (2n_q - 33)/48\pi$ where n_q is the number of quarks lighter than Q . With a coupling measured at $Q^2 = \mu_R^2$, the RGE facilitates predictions at other energy scales. The running of the SM gauge couplings is illustrated in Figure 1.3 (dashed lines). In particular the $SU(3)$ coupling shows a strong dependency on the energy scale, in contrast to those of $U(1)$ and $SU(2)$. In practice, experiments at the LHC therefore often treat the latter couplings as effectively constant within the range of energies and precision levels relevant to the analyses.

1.5 Beyond the Standard Model

1.4.2 Factorization and Parton Distribution Functions

While at high energy (low distance) scales perturbative QCD may be used to calculate interactions, at low energy (high distance) scales, the QCD coupling increases and perturbative treatments fail. For experiments in the LHC environment, an accurate description over the full Q^2 scale is important: parton interactions in proton beam collisions typically occur at high Q^2 , while the parton distribution inside the initial proton, as well as hadronization – only color singlets are allowed – are low Q^2 processes. Fortunately, the two regimes may be separated through the *factorization theorem* [22], such that the cross-section of $p_1 + p_2 \rightarrow f$ may be expressed as:

$$\sigma_{p_1 p_2 \rightarrow f} = \sum_{ij} \int f_i^{p_1}(x_1, \mu_F^2) f_j^{p_2}(x_2, \mu_F^2) \hat{\sigma}_{i+j \rightarrow f} dx_1 dx_2 \quad (1.28)$$

where $\hat{\sigma}_{i+j \rightarrow f}$ is the partonic cross-section, calculated with perturbative QCD, and $f_i^p(x_i, \mu_F^2)$ is the parton distribution function (PDF), ie. the probability density for parton i to carry a fraction x_i of the momentum of the original proton. The *factorization scale* μ_F defines the scale that separates the perturbative and non-perturbative regimes. The PDF corresponds to the non-perturbative dynamics involved in the proton-proton collision and cannot be calculated. That said, once a PDF is experimentally measured, QCD *does* describe its evolution via the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations:

$$\frac{df(x, Q^2)}{d\ln Q^2} = \frac{\alpha_s(Q^2)}{2\pi} \sum_j \int_x^1 \frac{dw}{w} P_{i \rightarrow j} \left(\frac{x}{w} \right) f_j(w, Q^2) \quad (1.29)$$

where $P_{i \rightarrow j}(z)$ is the splitting function, describing the probability of parton i to emit another parton j that carries a fraction $z \leq 1$ of its initial momentum. Examples of PDFs measured at $Q^2 = 10 \text{ GeV}^2$ and $Q^2 = 10 \text{ GeV}^4$ are shown in Figure 1.4.

1.5 Beyond the Standard Model

Once all free parameters of the theory are measured, the SM is fully defined and facilitates predictions for all fundamental particle interactions. The success of the SM theory is repeatedly demonstrated with experimental observations, which have already validated the predictions of SM cross-sections over more than 10 orders of magnitude [24]. Still, the SM does not provide answers to several fundamental

The Standard Model of Particle Physics

questions and the SM theory itself also comes with a range of peculiarities, as will be discussed below. This suggests the existence of new physics (NP) beyond the SM that is yet to be discovered, and motivates the NP searches presented in Chapters 5 and 6.

Gravity is the only fundamental force that is not accounted for in the SM. Its associated mass and distance scales ($\Lambda_{\text{Planck}} \simeq 10^{19}$ GeV) are many orders of magnitude larger than those of the electroweak and strong forces, and as such are not deemed relevant for SM physics processes. Furthermore, in contrast to the SM, general relativity is not a quantum theory. Quantizing gravity is not obvious, and generally results in non-renormalizable theories. Alternatively the unification of gravity and all other interactions could be achieved with string theories.

Dark matter and *dark energy* are known to make up roughly 21% and 74% of the mass/energy content of the universe, respectively, while only the remaining 5% is visible to us (mainly in the form of baryons). Evidence for Dark Matter stems from eg. gravitational effects, such as rotation curves of galaxies [25, 26], but the nature of it remains unclear. Many NP models, such as Supersymmetry, suggest the existence of weakly interacting massive particles (WIMPs), typically with masses in the range 10^2 - 10^3 GeV, neutral and stable. Given these masses are within reach of the LHC, both CMS and ATLAS aim to provide evidence these hypothetical particles, in addition to dedicated Dark Matter detectors.

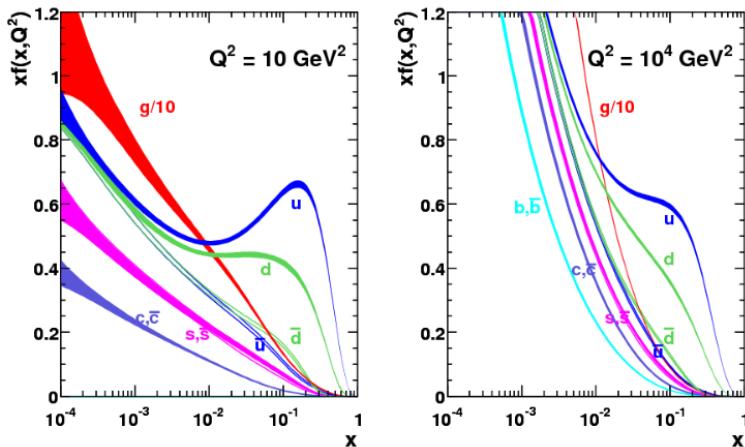


Figure 1.4: Next-to-leading order parton distribution function “MSTW 2008” for $Q^2=10$ GeV 2 and 10^4 GeV 2 [23]. The product $xf(x, Q^2)$ indicates the contribution of partons with momentum fraction x to the overall momentum distribution inside a hadron when probed at energy scale Q^2 .

1.5 Beyond the Standard Model

The *matter-antimatter asymmetry* problem refers to the observed asymmetry between baryons and antibaryons within the universe [27–29]. Since there are also no signs of local antimatter domains, which would balance out matter-dominated regions of the universe, it is inferred that some process in the early stages of the universe must have favored matter over antimatter. A recipe for such a process must satisfy the three Sakharov conditions [27]; These include *i*) the existence of a baryon (B) and lepton-number (L) violating process, *ii*) the occurrence of a period of thermal non-equilibrium, and *iii*) C- and CP-violation. Although the SM *can* accommodate processes that violate B and L via non-perturbative effects [30], the SM does not have the non-equilibrium condition. Also, while some CP-violation is observed in weak decays of quarks (associated to the phase factor of the CKM matrix), it is insufficient to explain the present matter-antimatter asymmetry. Furthermore, tight upper bounds on CP-violation in the strong interaction have been set via measurements of the electric dipole moment of the neutron (ie. the *strong CP problem*) [31, 32]. Some models of massive neutrinos accommodate new sources of CP violation in the lepton sector.

A *Grand Unified Theory* (GUT) considers an overarching symmetry group in which the SM gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$ is embedded, and the SM interactions are unified above its associated breaking scale m_X ($\gg m_Z$). An example of a GUT is the Minimal Supersymmetric Standard Model (MSSM) [19]; In this scenario the additional sparticles modify the β functions in the renormalization group equations (Equation 1.26), such that the couplings have equal strength at $m_X \simeq 10^{16}$ GeV, as shown in Figure 1.3 (solid lines).

The *hierarchy problem* is related to the mass of the scalar Higgs boson; It can be shown that the tree-level Higgs mass receives higher order corrections proportional to Λ^2 , where Λ is the next higher energy scale in the theory. If the next energy scale is GUT ($\Lambda_{\text{GUT}} \simeq 10^{16}$ GeV) or gravity ($\Lambda_{\text{Planck}} \simeq 10^{19}$ GeV), then the corrections to the bare Higgs mass would be more than 30 orders of magnitude larger than its observed value of $\simeq 125$ GeV. Only a fine-tuning of astronomical precision would cause the corrections to cancel down to the level of the observed Higgs mass. Although this strictly speaking is not a problem of consistency of the theory, it *is* a problem of naturalness. In particular Supersymmetry is an attractive solution to the hierarchy problem, as the loop corrections introduced by the new sparticles have opposite sign with respect to their SM counterparts, hence making the cancellations by construction exact.

Chapter 2

Supersymmetry

Several fundamental questions, such as the nature of Dark Matter and the hierarchy problem, request new physics phenomena beyond the SM (BSM). Among the many BSM theories, Supersymmetry [8–12] is arguably the most remarkable one, predicting the existence of a new *superpartner* for each of the SM particles.

This chapter first motivates the existence of Supersymmetry at the energy scale probed by the LHC and provides an unapologetically short overview of the theoretical framework. From Section 2.3 onwards, the focus shifts towards the experimental aspects, involving simplified Supersymmetry models, their underlying assumptions and general search strategies. Furthermore, several Supersymmetry scenarios with compressed mass-spectra and/or long-lived particles are introduced, which yield unique and hard-to-probe final state signatures. These scenarios are targeted in the search presented in Chapter 5. The discussions are largely based on Refs. [13, 16, 19, 33].

Supersymmetry

2.1 Motivation for Supersymmetry at the TeV-Scale

In Section 1.5, several shortcomings of the SM where mentioned, and new theories beyond the SM (BSM) are continuously being explored in order to answer to those. Many such theories predict the existence of new phenomena at the electroweak (TeV) scale, within reach of the physics potential of the LHC. Supersymmetry (SUSY) has been around since the early 1970's and is arguably one of the most remarkable BSM theories, suggesting new particle states at the TeV scale. It consists of a new symmetry that relates bosonic and fermionic degrees of freedom via the Supersymmetric transformation operator Q :

$$Q|\text{Boson}\rangle = |\text{Fermion}\rangle, \quad Q|\text{Fermion}\rangle = |\text{Boson}\rangle \quad (2.1)$$

The SM particle content is expanded such that each SM particle gains a *superpartner* that differs in spin by half a unit. The associated SM and superpartner fields are organized into *supermultiplets*, which are a function of *superspace*, a generalization of spacetime that adds new fermionic coordinates. A few more details on the structure of SUSY models will be given in Section 2.2, but the main implications and motivations for TeV-scale Supersymmetric new physics phenomena are already noted in the following.

Hierarchy problem – As mentioned in Section 1.5, the tree-level Higgs mass is subject to higher order corrections corresponding to loop contributions to the Higgs propagator. For loops containing a fermion f with mass m_f , the corrections are of the form:

$$\Delta m_H^2 = -\frac{|\lambda_f|^2}{8\pi^2} \Lambda^2 + \dots \quad (2.2)$$

where λ_f is the coupling of the 3-point interaction between f and the Higgs, and Λ is the next higher energy scale at which new physics appears. Similarly, loops containing a scalar S with mass m_S and 4-point coupling to the Higgs λ_S , contribute as:

$$\Delta m_H^2 = +\frac{\lambda_S}{16\pi^2} \Lambda^2 + \dots \quad (2.3)$$

An example of the corresponding loop diagrams is shown in Figure 2.1. In case $\Lambda = \Lambda_{\text{Planck}} \simeq 10^{19}$ GeV, the quantum corrections to the Higgs mass would be more than 30 orders of magnitude larger than the observed value of 125 GeV. The quadratic dependency on the mass of the heaviest particle that the Higgs couples to is clearly a problem of electroweak scale instability. The tendency to solve this problem by simply letting Λ be small often results in problems with unitarity or

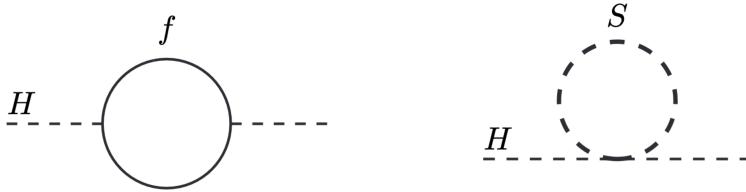


Figure 2.1: Diagrams corresponding to one-loop correction to the Higgs mass due to a fermion (left) or scalar (right).

causality [34]. Alternatively, one could consider a scenario where none of the new high-mass particles couple to the Higgs field, although this is a rather puzzling and unsatisfying assumption. It therefore appears that the observed Higgs mass is the result of cancellations between the various contributions to Δm_H^2 down to astronomical precision. That said, Supersymmetry implies that this cancellation is in fact not a perplexing coincidence, but rather an exact result of the Supersymmetric relation between fermions and bosons; Each SM fermion would be accompanied by two real scalars – or equivalently one complex scalar – organized into a single supermultiplet with $|\lambda_f|^2 = \lambda_S$. In this case, the first terms in Equations 2.2 and 2.3 cancel, meaning that the new particle states protect the Higgs mass from the quadratic divergencies. Supersymmetry hence provides a natural solution to the hierarchy problem. It should be noted, however, that the most natural realization of Supersymmetry implies that the SM particles and their superpartners are mass-degenerate. The absence of observed superpartners excludes this possibility, and as such Supersymmetry must be a broken symmetry. It may nevertheless still be a solution to the hierarchy problem if the Supersymmetry breaking is soft, accepting a small amount of finetuning in the cancellation of the Higgs mass corrections. This corresponds to Supersymmetry breaking mass parameters (discussed below, Equation 2.7) of maximally a few TeV. Naturalness imposes constraints in particular on the masses of higgsinos, top squarks (often assumed the lightest of the squarks) and gluinos [35–41]. While it is difficult to pose exact upper bounds, several studies, such as in Refs. [42, 43], suggest naturalness can be maintained with higgsino masses of $\lesssim 300$ GeV, top squark masses of $\lesssim 1$ TeV and gluino masses of $\lesssim 2$ TeV.

Dark Matter – Conservation of $B-L$ is an important property that contributes to the observed proton stability and suppression of neutrino masses. While in the SM theory this is an accidental symmetry, in generic Supersymmetry this is not the case. When constructing the Supersymmetric Lagrangian, $B-L$ conservation is therefore often *imposed*. Consequently, such a Supersymmetry model gains

Supersymmetry

the property of R-parity invariance, with R defined as:

$$R = (-1)^{3(B-L)+2S} \quad (2.4)$$

It can be shown that SM particles have even R-parity and their superpartners have odd R-parity. As such, R-parity conservation has several important implications for Supersymmetric phenomenology; In collider experiments, such as the LHC, the initial state particles are always R-parity even. If Supersymmetric particles are produced in the collision, they *must* be produced in pairs to maintain R-parity conservation. Secondly, a given R-parity odd Supersymmetric particle can never decay to only R-parity even (SM) particles. As such, the lightest Supersymmetric particle (LSP) *must* be stable. An electrically and color neutral LSP would interact only weakly with SM matter, and it is shown that such a WIMP with mass around the TeV scale would be consistent with the observed Dark Matter density [44]. The LSP can therefore serve as a suitable Dark Matter candidate. If produced at experiments like CMS, the LSP would behave similar to a SM neutrino, yielding no distinct detector signature except for a contribution to the total missing transverse energy in the final state. Lastly, $B-L$ conservation *can* be omitted from the theory, resulting in models with R-parity violation. In these cases, the LSP is unstable and is no longer a viable candidate for Dark Matter. Additionally, it would generate new B and L violating processes that could yield a proton decay rate larger than the observed bounds.

Grand unification – As already alluded to in Section 1.5, a minimal Supersymmetric extension of the SM would introduce new particle states that modify the β functions in the renormalization group equations. With the particle content of the MSSM, starting below a few TeV, an approximate unification of the SM gauge couplings at high energy ($m_X \simeq 10^{16}$ GeV) can be achieved [45]. Although the importance of grand unification may be questioned, the fact that it follows as a consequence of TeV scale Supersymmetry is in itself a striking achievement.

Recent experimental anomalies – Several recent low-energy measurements that are sensitive to Supersymmetric loop effects have shown a deviation from the theory expectation. Although one should avoid interpreting such deviations as evidence for Supersymmetry, the compatibility with the existence of Supersymmetric particles at the electroweak scale is noteworthy. For example, the Fermilab Muon g-2 collaboration recently published measurements of the muon anomalous magnetic moment, showing deviations with the theory prediction of up to 5σ [46, 47]. Although recent lattice calculations provide an ambiguous the-

2.2 The Minimal Supersymmetric Standard Model

ory prediction [48, 49], it has been shown that there is viable parameter space for Supersymmetric realizations with sizable contributions to the muon g-2 and light Supersymmetric particles (eg. smuon masses below $\simeq 300$ GeV) that are compatible with current LHC constraints, and would answer to the apparent muon g-2 anomaly [50]. Another deviation corresponds to the recent W boson mass measurement by the CDF collaboration, which claimed record precision and 7σ tension with the SM expectation [51]. While satisfying the relevant experimental constraints on the muon g-2, Dark Matter relic density and LHC searches, TeV scale Supersymmetric particles *could* explain the observed shift in W boson mass (with $m_{\tilde{t}_1}$ up to $\simeq 1.5$ TeV) [52]. That said, continued effort on both experimental and theoretical sides are needed to conclusively confirm or reject these anomalies.

2.2 The Minimal Supersymmetric Standard Model

This section gives an overview of the Minimal Supersymmetric Standard Model (MSSM), which extends the SM with a minimally enlarged Higgs sector and introduces superpartners for each SM field. More complex Supersymmetric extensions clearly do exist, such as the Next-to-minimal Supersymmetric Standard Model (NMSSM) [53] or generalized NMSSM (GNMSSM) [54], which contain extra structure to explain the observed neutrino masses or address theoretical or experimental tensions. However, these are outside the scope of this thesis and will be ignored in the following.

2.2.1 Theoretical Structure

The MSSM is constructed by expanding the Higgs sector of the SM with a second complex Higgs doublet and adding superpartners to each SM field. The SM gauge bosons and their fermionic superpartners (gauginos) together form gauge (or vector) multiplets. The leptons and quarks together with their scalar superpartners (sleptons and squarks, respectively), as well as the two complex Higgs doublets with their fermionic superpartners (higgsinos) form chiral multiplets. The MSSM field content is summarized in Table 2.1.

The Supersymmetric Lagrangian depends on three functions of the superfields Φ : the Kähler potential K , gauge kinetic function f_{ab} and superpotential W .

Supersymmetry

Superfield	Bosonic field	Fermionic field	Representations		Quantum numbers		
			$SU(3)_c$	$SU(2)_L$	t_L^3	y	$q = t_L^3 + y$
\widehat{Q}	$\tilde{q}_{mL} = \begin{pmatrix} \tilde{u}_m \\ \tilde{d}_m \end{pmatrix}_L$	$q_{mL} = \begin{pmatrix} u_m \\ d_m \end{pmatrix}_L$	3	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$+\frac{1}{6}$	$\begin{pmatrix} +\frac{2}{3} \\ -\frac{1}{3} \end{pmatrix}$
\widehat{U}^c	\tilde{u}_{mR}^-	u_{mR}	3	1	0	$+\frac{2}{3}$	$+\frac{2}{3}$
\widehat{D}^c	\tilde{d}_{mR}^-	d_{mR}	3	1	0	$-\frac{1}{3}$	$-\frac{1}{3}$
\widehat{L}	$\tilde{\ell}_{mL} = \begin{pmatrix} \tilde{\nu}_m \\ \tilde{e}_m \end{pmatrix}_L$	$\ell_{mL} = \begin{pmatrix} \nu_m \\ e_m \end{pmatrix}_L$	1	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$-\frac{1}{2}$	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$
\widehat{E}^c	\tilde{e}_{mR}^-	e_{mR}	1	1	0	-1	-1
\widehat{V}_8	G	G	8	1	0	0	0
\widehat{V}	$W = \begin{pmatrix} W^+ \\ W^- \\ W^0 \end{pmatrix}$	$\widetilde{W} = \begin{pmatrix} \widetilde{W}^+ \\ \widetilde{W}^- \\ \widetilde{W}^0 \end{pmatrix}$	1	3	$\begin{pmatrix} +1 \\ -1 \\ 0 \end{pmatrix}$	0	$\begin{pmatrix} +1 \\ -1 \\ 0 \end{pmatrix}$
\widehat{V}'	B	\widetilde{B}	1	1	0	0	0
\widehat{H}_u	$H_u = \begin{pmatrix} H_u^+ \\ H_u^0 \end{pmatrix}$	$\widetilde{H}_u = \begin{pmatrix} \widetilde{H}_u^+ \\ \widetilde{H}_u^0 \end{pmatrix}$	1	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$+\frac{1}{2}$	$\begin{pmatrix} +1 \\ 0 \end{pmatrix}$
\widehat{H}_d	$H_d = \begin{pmatrix} H_d^0 \\ H_d^- \end{pmatrix}$	$\widetilde{H}_d = \begin{pmatrix} \widetilde{H}_d^0 \\ \widetilde{H}_d^- \end{pmatrix}$	1	2	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$-\frac{1}{2}$	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$

Table 2.1: Supersymmetric version of Table 1.1, showing the fields in the MSSM and their representations under $SU(3)_c$ and $SU(2)_L$, as well as the quantum numbers associated to $SU(2)_L$, $U(1)_Y$ and $U(1)_Q$ (which emerges after $SU(2)_L \times U(1)_Y \rightarrow U(1)_Q$ electroweak symmetry breaking). Superpartners of the SM fields are indicated with a tilde. The boson fields and their corresponding fermionic partners differ in spin by half a unit and are contained in superfields, denoted with hatted symbols. Bold numbers represent the dimension of the representations: **1**, **2**, **3** and **8** denote a singlet, doublet, triplet and octet, respectively. The superscript ⁰ for the (s)fermion weak eigenstates has been omitted to avoid confusion with the neutral components of the Higgs doublets, while subscript _{*m*} ($m \in \{1, 2, 3\}$) corresponds to the (s)fermion families.

2.2 The Minimal Supersymmetric Standard Model

An example of a general non-renormalizable¹ gauge-invariant Lagrangian has the form:

$$\mathcal{L} = \left[K(\Phi_i, \tilde{\Phi}^{*j}) \right]_D + \left(\left[\frac{1}{4} f_{ab}(\Phi_i) \widehat{\mathcal{W}}^{a\alpha} \widehat{\mathcal{W}}^b_\alpha + W(\Phi_i) \right]_F + c.c. \right) \quad (2.5)$$

where $\widehat{\mathcal{W}}$ denotes normalized field-strength superfields. The Kähler potential K and gauge kinetic function in their simplest form yield standard kinetic energy terms for all fields. The superpotential in the MSSM, which generates Supersymmetric Yukawa vertices and the Higgs(ino) masses, is written as:

$$W = \lambda_d \widehat{H}_d \widehat{Q} \widehat{D}^c - \lambda_u \widehat{H}_u \widehat{Q} \widehat{U}^c + \lambda_e \widehat{H}_d \widehat{L} \widehat{E}^c + \mu \widehat{H}_u \widehat{H}_d \quad (2.6)$$

Interaction terms are then added, which adhere to the $SU(3) \times SU(2) \times U(1)$ gauge group of the SM and maintain $B - L$ invariance. Finally, general soft-Supersymmetry breaking terms are needed for a viable theory given the constraints from experiments. These may be written as:

$$\begin{aligned} \mathcal{L}_{soft} = & -\frac{1}{2} \left(M_1 \tilde{B} \tilde{B} + M_2 \tilde{W} \tilde{W} + M_3 \tilde{g} \tilde{g} + c.c. \right) \\ & - \left(\tilde{U}^c A_u \tilde{Q} H_u - \tilde{D} A_d \tilde{Q} H_d - \tilde{E}^c A_e \tilde{L} H_d + c.c. \right) \\ & - \tilde{Q}^\dagger M_{\tilde{Q}}^2 \tilde{Q} - \tilde{L}^\dagger M_{\tilde{L}}^2 \tilde{L} - \tilde{U}^c M_{\tilde{U}^c}^2 \tilde{U}^{c\dagger} - \tilde{D}^c M_{\tilde{D}^c}^2 \tilde{D}^{c\dagger} - \tilde{E}^c M_{\tilde{E}^c}^2 \tilde{E}^{c\dagger} \\ & - m_{H_u}^2 |H_u|^2 - m_{H_d}^2 |H_d|^2 - (b H_u H_d + c.c.) \end{aligned} \quad (2.7)$$

Here, the first line consists of the bino, wino and gluino mass-terms. The second line contains the tri-scalar couplings, with $A_{u,d,e}$ representing complex 3×3 matrices (for 3 families). The third line corresponds to the squark and slepton mass terms, with $M_{\tilde{Q}, \tilde{L}, \tilde{U}^c, \tilde{D}^c, \tilde{E}^c}$ denoting hermitian 3×3 matrices (again for 3 families). Finally, the last line contains Supersymmetry-breaking contributions to the Higgs potential.

The soft-Supersymmetry breaking Lagrangian of Equation 2.7 introduces the majority of free parameters of the MSSM and therefore adds a large arbitrariness to the theory. Furthermore, the fundamental mechanism that achieves soft-Supersymmetry breaking is probably the most unclear feature within the frame-

¹Non-renormalizable Supersymmetric Lagrangians are hard to avoid when eg. accounting for the effect of very heavy states on low-energy interactions or incorporating gravity in the theory. However, those non-renormalizable interactions can often be ignored, since they are suppressed by inverse orders of the heavy mass-scale.

Supersymmetry

work of generic Supersymmetry. It might be broken spontaneously, in a way similar to the electroweak symmetry breaking in the SM, which would imply the existence of a massless goldstino fermion in case Supersymmetry is a global symmetry. Conversely, if Supersymmetry is a local symmetry, then the goldstino would be absorbed into the gravitino, incorporating gravity in the theory (known as supergravity), via the *super-Higgs mechanism* [55]. That said, observational limits suggest that Supersymmetry breaking probably does not occurs as a consequence of MSSM particle interactions alone. Instead there may be a *hidden sector*² in which Supersymmetry is broken, which is then propagated to the visible MSSM sector via a mediation mechanism. Among the many mediation mechanisms, the two most commonly mentioned are gravity-mediated and gauge-mediated Supersymmetry breaking. In the former, gravity is the messenger connecting the visible MSSM and hidden Supersymmetry breaking sectors via effects that are suppressed by inverse powers of the Planck mass. In the gauge-mediated Supersymmetry breaking (GMSB) mechanism [56], the breaking is transmitted by messenger fields, being new chiral multiplets with non-trivial SM quantum numbers. In GMSB models the gravitino would be the LSP and a potential candidate for Dark Matter.

It should be noted that soft-Supersymmetry breaking is not strictly needed *if* one is prepared to give up on Supersymmetry as a natural solution to the hierarchy problem. In that case, Supersymmetry could be decoupled from the origin of the electroweak scale and Supersymmetry breaking may appear at vastly larger scales, implying that fine-tuning of the Higgs mass corrections again becomes an issue. It could, however, still yield viable candidates for Dark Matter and achieve grand unification. In models of *split-Supersymmetry* the squarks and sleptons would be too heavy to be produced at the LHC (masses up to 10^{10} TeV), but there would still be fermionic superpartners of the gauge and Higgs bosons that have masses on the TeV scale or even below [57–59]. A “compromise” is introduced by *mini-split-Supersymmetry* models, that reduce the mass-splitting between fermionic and bosonic superpartners by forcing scalar masses to be below 10^5 TeV, hereby facilitating more easily the 125 GeV Higgs [60, 61].

Finally, the extension of the scalar Higgs sector with a second complex Higgs doublet implies a modified description of $SU(2)_L \times U(1)_Y \rightarrow U(1)_Q$ electroweak symmetry breaking. The Higgs potential in the MSSM may be written as:

²The existence of hidden sectors is a concept not unique to Supersymmetry.

2.2 The Minimal Supersymmetric Standard Model

$$V^{\text{MSSM}} = (|\mu|^2 + m_{H_u}^2)|H_u^0|^2 + (|\mu|^2 + m_{H_d}^2)|H_d^0|^2 - (bH_u^0H_d^0 + c.c.) + \frac{1}{8}(g^2 + g'^2)(|H_u^0|^2 + |H_d^0|^2)^2 \quad (2.8)$$

It can be shown that when requiring the potential to be bounded from below as $2b < 2|\mu|^2 + m_{H_u}^2 + m_{H_d}^2$, and choosing $b^2 > (|\mu|^2 + m_{H_u}^2)(|\mu|^2 + m_{H_d}^2)$, electroweak symmetry breaking occurs as consequence of non-zero vacuum expectation values:

$$\nu_u = \langle H_u^0 \rangle, \quad \nu_d = \langle H_d^0 \rangle, \quad \tan\beta \equiv \nu_u/\nu_d \quad (2.9)$$

2.2.2 Particle Content

Regardless of the admittedly large degree of uncertainty concerning the Supersymmetry breaking mechanism, the physical Supersymmetric particles that are sought after at the LHC are mass-eigenstates of the fields presented in Table 2.1.

Charginos and neutralinos – The winos, bino and higgsinos mix due to SM electroweak symmetry breaking effects; The four mass-eigenstate charginos ($\tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm$) are obtained through linear combinations of the two charged winos (\tilde{W}^\pm) and two charged higgsinos (\tilde{H}_u^+ and \tilde{H}_d^-), which is described via the following 2×2 mass-mixing matrix:

$$M_C = \begin{pmatrix} M_2 & \frac{g\nu_u}{\sqrt{2}} \\ \frac{g\nu_d}{\sqrt{2}} & \mu \end{pmatrix} \quad (2.10)$$

Diagonalization with unitary matrices U and V yields the four chargino masses:

$$U^* M_C V^{-1} = \text{diag}(M_{\tilde{\chi}_1^\pm}, M_{\tilde{\chi}_2^\pm}) \quad (2.11)$$

The mixing of the bino (\tilde{B}), neutral wino (\tilde{W}^0) and two neutral higgsinos (\tilde{H}_u^0 and \tilde{H}_d^0) yields the four mass-eigenstate neutralinos (χ_i^0 with $i \in \{1, 2, 3, 4\}$), as characterized by the 4×4 mass-matrix:

$$M_N = \begin{pmatrix} M_1 & 0 & -\frac{g'\nu_d}{2} & \frac{g'\nu_u}{2} \\ 0 & M_2 & \frac{g\nu_d}{2} & -\frac{g\nu_u}{2} \\ -\frac{g'\nu_d}{2} & \frac{g\nu_d}{2} & 0 & -\mu \\ \frac{g'\nu_u}{2} & -\frac{g\nu_u}{2} & -\mu & 0 \end{pmatrix} \quad (2.12)$$

The four neutralino masses are similarly obtained, using unitary matrix W :

$$W^T M_N W = \text{diag}(M_{\tilde{\chi}_1^0}, M_{\tilde{\chi}_2^0}, M_{\tilde{\chi}_3^0}, M_{\tilde{\chi}_4^0}) \quad (2.13)$$

Gluinos – The gluino is the only $SU(3)_c$ octet, hence cannot mix with other particles. The gauge and mass-eigenstates are therefore the same.

Squarks and sleptons – The mass-eigenstates of the squarks and sleptons are obtained by diagonalizing three 6×6 matrices: one for up-type squarks, one for down-type squarks and one for charged sleptons, each in the basis of three sfermion families and two chiralities. In general, however, the mixing angles for the first and second families are predicted to be negligible, leaving the possibility of $\tilde{f}_L - \tilde{f}_R$ mixing only in the third family. This conveniently reduces the exercise to diagonalizing 2×2 matrices for the squarks and charged sleptons. In the absence of a right-handed sneutrino in the MSSM, the sneutrino masses are obtained by diagonalizing one 3×3 matrix. Nevertheless, the explicit expressions for the sfermion masses are complicated when considering eg. contributions beyond the tree-level or RGEs, and will not be given here. Instead, more details can be found in eg. Ref. [19].

Higgs bosons – The two complex $SU(2)_L$ Higgs doublets together have 8 degrees of freedom. When electroweak symmetry is broken, three are absorbed into the massive Z^0 and W^\pm bosons. The remaining degrees of freedom yield 5 scalar mass-eigenstates: two neutral CP-even scalars h^0 (the lightest and often assumed to be the observed 125 GeV scalar) and H^0 , one neutral CP-odd scalar A^0 and two charged scalars H^\pm . At tree-level, their masses can be expressed as:

$$m_{A^0}^2 = 2b/\sin(2\beta) = 2|\mu|^2 + m_{H_u}^2 + m_{H_d}^2 \quad (2.14)$$

$$m_{h^0, H^0}^2 = \frac{1}{2} \left(m_{A^0}^2 + m_Z^2 \mp \sqrt{(m_{A^0}^2 - m_Z^2)^2 + 4m_Z^2 m_{A^0}^2 \sin(2\beta)} \right) \quad (2.15)$$

$$m_{H^\pm}^2 = m_{A^0}^2 + m_W^2 \quad (2.16)$$

2.3 Searches for Supersymmetry at the LHC

A careful count reveals that the MSSM contains 124 free parameters (masses, mixing angles, CP-violating phases, etc.), most of which are introduced by the soft-Supersymmetry breaking terms (Equation 2.7). Experimental observations clearly constrain the MSSM parameter freedom, but the large dimensionality of the parameter space makes it difficult to interpret the observations in fully generic

2.3.1 Simplified Models

One way to reduce the number of free parameters of the MSSM is to assume a simple structure in the soft-Supersymmetry breaking terms at a high energy scale Λ_X , typically that of grand unification. In models of minimal supergravity (mSUGRA) or constrained MSSM (cMSSM) [9, 62], one imposes that the scalar squared-masses and couplings are diagonal and universal: $M_{\tilde{Q}, \tilde{U}^c, \tilde{D}^c, \tilde{L}, \tilde{E}^c}^2(\Lambda_X) = m_0^2 \mathbf{1}$ and $A_{u,d,e} = A_0 \mathbf{1}$, with $\mathbf{1}$ denoting the 3×3 unity matrix. Furthermore, grand unification predicts that at Λ_X the gaugino mass-parameters are the same: $M_{1,2,3}(\Lambda_X) = m_{1/2}$. These relations can then be used as initial conditions for the renormalization group equations to derive expressions for Supersymmetric parameters at energy scales that are within reach of experiments. In fact, with the above assumptions it turns out that complete mass-spectrum of the MSSM particles is governed by only 5 parameters: m_0 , $m_{1/2}$, A_0 , and the values for μ and $B = b/\mu$ at Λ_X .

Alternatively, experiments at the LHC often interpret their search results using *Simplified Model Spectra* (SMS) [63, 64]. These consist of an effective Lagrangian where only the lightest (typically two) Supersymmetric particles are included, while the others are assumed to be decoupled and too heavy to be produced with proton collisions at LHC energies. This allows one to probe the model space in terms of only the properties of the LSP and next-to-lightest Supersymmetric particle (NLSP), which is additionally motivated by the limited computational power available to produce the Monte-Carlo events for each signal hypothesis, and facilitates visualization of the experimental results with bounds on the $(m_{\text{NLSP}}, m_{\text{LSP}})$ mass-plane. In the computation of the signal cross-sections it is generally assumed that the MSSM mixing of gauge-eigenstates into mass-eigenstates occurs such that the mass-eigenstates are (nearly) pure in one of the gauge-eigenstates. This implies an assumption on the underlying MSSM mass-parameters; For example, if $|M_1| \ll |M_2|, \mu$, then the lightest neutralino ($\tilde{\chi}_1^0$) would be bino-like. If $|M_2| \ll |M_1|, \mu$, then $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^0$ would be a triplet of nearly mass-degenerate pure winos. If $|\mu| \ll |M_1|, |M_2|$, then $\tilde{\chi}_1^\pm$, $\tilde{\chi}_1^0$ and $\tilde{\chi}_2^0$ would be nearly mass-degenerate higgsinos. An overview of cross-sections of pair-produced sparticles in SMS for proton-proton collisions at $\sqrt{s} = 13$ TeV is shown in Figure 2.2.

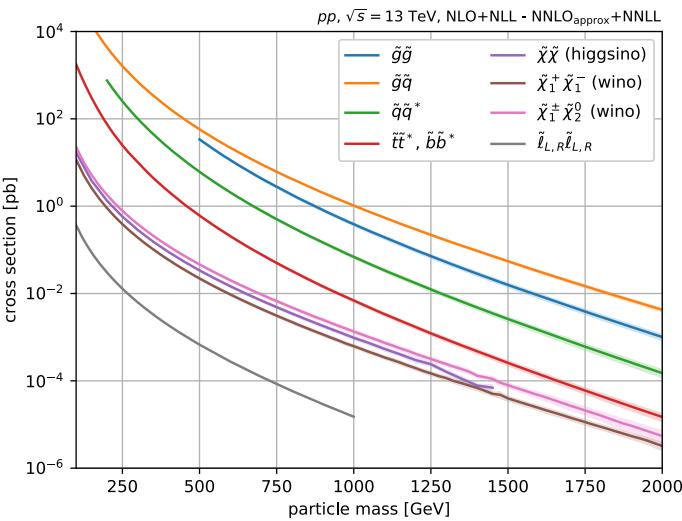


Figure 2.2: Cross-sections of pair-produced Supersymmetric particles in SMS models for proton-proton collisions at a center-of-mass energy of 13 TeV [65].

The scenarios discussed above rely on assumptions that have to be tested experimentally and clearly might be false. A more model-independent way of probing the MSSM model space can be done by interpreting experimental results in terms of the *phenomenological MSSM* (pMSSM) [66–68]. This model consists of only 19 free parameters and captures most of the electroweak scale phenomenology of the MSSM, based on the following assumptions: *i*) R-parity conservation, *ii*) no new sources of CP violation beyond those in the CKM matrix, *iii*) minimal flavor violation [69], *iv*) the first two sfermion generations are mass-generate and have couplings too small to be experimentally relevant, and *v*) $\tilde{\chi}_1^0$ is the LSP.

2.3.2 Experimental Search Strategies

The phenomenology of the new particle content clearly depends on the targeted realizations and underlying assumptions. As mentioned before, experiments typically search for the lightest Supersymmetric particles in SMS models, which are assumed light enough to be produced at the LHC, while the other sparticles are decoupled. Experimental signatures of the signal hypotheses are then to be distinguished from SM background. At the hadronic environment of the LHC, the SM background is dominated by QCD multijet events, while high parton-parton col-

lision energies also yield significant contributions from (associated) vector boson production and top quark physics.

In R-parity conserving Supersymmetry models, the sparticles are pair-produced and decay via emission of their SM partners to a stable LSP, which escapes the experiment undetected. Final states are therefore composed of the visible SM decay products as well as missing transverse energy (p_T^{miss}) induced by the LSP. As such, in most searches the p_T^{miss} serves as an important handle to distinguish signal from SM background. Searches for squarks and gluinos typically involve final states with significant p_T^{miss} as well as high energy jets from heavy sparticle decays, leading to high scalar sums of jet transverse momenta (H_T). Both the H_T and jet properties, such as jet substructure [70], can be exploited to reject SM backgrounds. Conversely, searches for charginos and neutralinos (collectively called *electroweakinos*) or sleptons typical involve final states with notable p_T^{miss} , accompanied with high energy leptons from SM W/Z boson or direct slepton decays. In these searches, properties of the leptons, such isolation and identification criteria, are important handles to reduce hadronic background.

Other selection variables have been defined to distinguish Supersymmetric event topologies from SM background. For collision events with Supersymmetric particles, the effective mass, defined as $m_{\text{eff}} = H_T + p_T^{\text{miss}}$, typically peaks around the mass-scale of the sparticles, while for SM background it rapidly drops and has only small tails at high m_{eff} . More examples that have proven useful include: $\alpha_T = p_T^{j2}$ [71], the stransverse mass M_{T2} [72], the cotransverse mass M_{CT} [73] and razor variables [74].

Lastly, in many searches the analysis regions are often defined by selections on kinematic variables such as mentioned above. Signal regions are designed to drive the sensitivity, typically for a range of signal hypotheses, while control regions may constrain the modeling of important background processes using the observed data events. After *cutting*, the analysis sensitivity is then gauged by simply *counting* the final expected signal and background events (hence referred to as *cut-and-count* analyses). Moreover, the use of machine learning techniques, such as boosted decision trees and neural networks, has become more popular in recent years, allowing one to perform more advanced object and event selections, exploiting complex correlations between multiple features. This typically leads to improved analysis sensitivity, as is shown in eg. Ref. [75].

Supersymmetry

2.3.3 Run 2 Results from the CMS Experiment

During the LHC Run 2 data-taking period (2016-2018), the CMS experiment collected approximately 137 fb^{-1} of proton-proton collision data at a center-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$. This data-set, together with state-of-the-art analysis techniques, enabled the extensive search program of the CMS Collaboration to reach unprecedented sensitivity to Supersymmetry models, leading to stringent bounds on the sparticle masses and production cross-sections [77–98]. To provide a sense of the physics reach of CMS after Run 2, a non-exhaustive summary is shown below, corresponding to R-parity conserving SMS models. A more complete review of the constraints on the MSSM parameter space is given in eg. Ref. [16].

Gluinos and squarks – Among the superpartners, colored sparticles have the highest cross-sections at the LHC, due to their production via the strong interaction. As such, the most stringent mass-constraints are obtained for gluinos and squarks. Exclusion limits for pair-produced gluinos are shown in Figure 2.3.

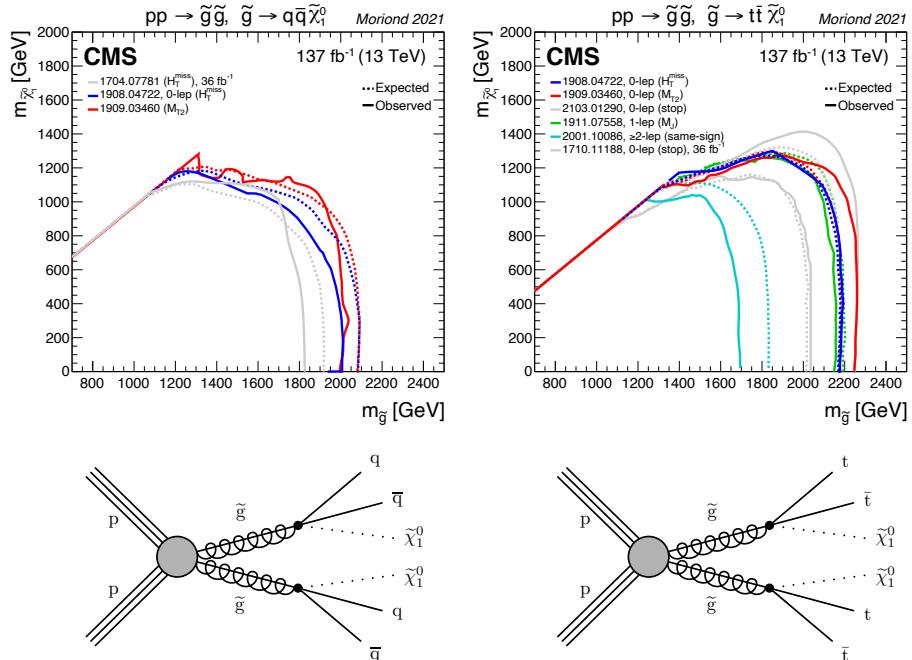


Figure 2.3: Exclusion limits for pair-produced gluinos in SMS models. The corresponding decay topologies are indicated with the diagrams below, where q represents a quark of the first two families. Figures from [76], with individual publications indicated in the legends.

2.3 Searches for Supersymmetry at the LHC

Each gluino is assumed to decay to a pair of quarks and an LSP neutralino. If the quarks are of the first two families, the final states consist of multiple jets and p_T^{miss} . In this scenario, gluinos are excluded up to 2 TeV for light neutralinos. For heavier neutralinos, the visible jets become less energetic, leading to reduced sensitivity. Conversely, if the gluino decays to top quarks, the final states consist of two $\tilde{\chi}_1^0$ plus the decay product of four top quarks, leading to rich signatures involving isolated leptons, b jets, light jets and p_T^{miss} . In this case, gluinos up to 2.2 TeV are excluded for light $\tilde{\chi}_1^0$. The mass-bounds for squarks are shown in Figure 2.4. Squarks of the first two families are generally assumed to be mass-degenerate, implying a production cross-section that is 8 times higher (4 light squarks times 2 chiralities) than that of a single squark. The pair-produced light squarks each are assumed to decay as $\tilde{q} \rightarrow q \tilde{\chi}_1^0$, and limits are set on squark masses up to approximately 1.8 TeV. Without the assumption of mass-degenerate light squarks, bounds reach only 1.3 TeV. For pair-produced top squarks, with $\tilde{t} \rightarrow t \tilde{\chi}_1^0$ and depending on the $\tilde{\chi}_1^0$ mass, top squark masses up to 1.3 TeV are excluded as well.

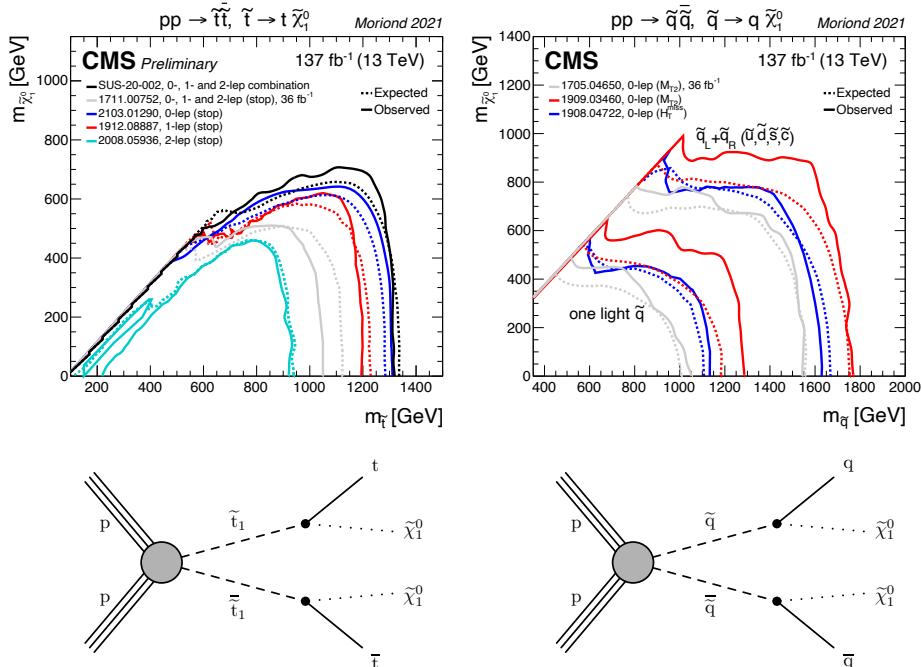


Figure 2.4: Exclusion limits for pair-produced squarks in SMS models. The corresponding decay topologies are indicated with the diagrams below, where q (\tilde{q}) represents a quark (squark) of the first two families. Figures from [76], with individual publications indicated in the legends.

Supersymmetry

Charginos and neutralinos – Figure 2.5 shows the exclusion limits for scenarios with wino-like mass-degenerate $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$. Since sleptons are assumed to be much heavier, chargino decays may occur only as $\tilde{\chi}_1^\pm \rightarrow W^\pm \tilde{\chi}_1^0$, while the NLSP neutralino decays as $\tilde{\chi}_2^0 \rightarrow Z \tilde{\chi}_1^0$ or $\tilde{\chi}_2^0 \rightarrow h \tilde{\chi}_1^0$. In the WH topology, searches are most sensitive if they exploit the $h \rightarrow bb$ decay channel. All hadronic final states allow exclusion of $\tilde{\chi}_1^\pm$ masses up to roughly 1 TeV for light $\tilde{\chi}_1^0$, while the presence of an isolated lepton from W^\pm decays improves sensitivity for heavier $\tilde{\chi}_1^0$. In the WZ topology, searches typically require the presence of 2 or 3 leptons in the final state in addition to p_T^{miss} . Exclusion limits reach $\tilde{\chi}_1^\pm$ masses up to 750 GeV for light $\tilde{\chi}_1^0$. Dedicated analysis strategies are required to target the phase-space where the mass-splitting between the NLSP and LSP is low, in which case the visible decay products are very soft and hard to distinguish from the overwhelming QCD background. This will be thoroughly discussed in Section 2.4 and Chapter 5. Wino-like $\tilde{\chi}_1^\pm$ masses up 275 GeV are excluded for a mass-splitting of 10 GeV. The experimental bounds on the production of higgsino-like electroweakinos are

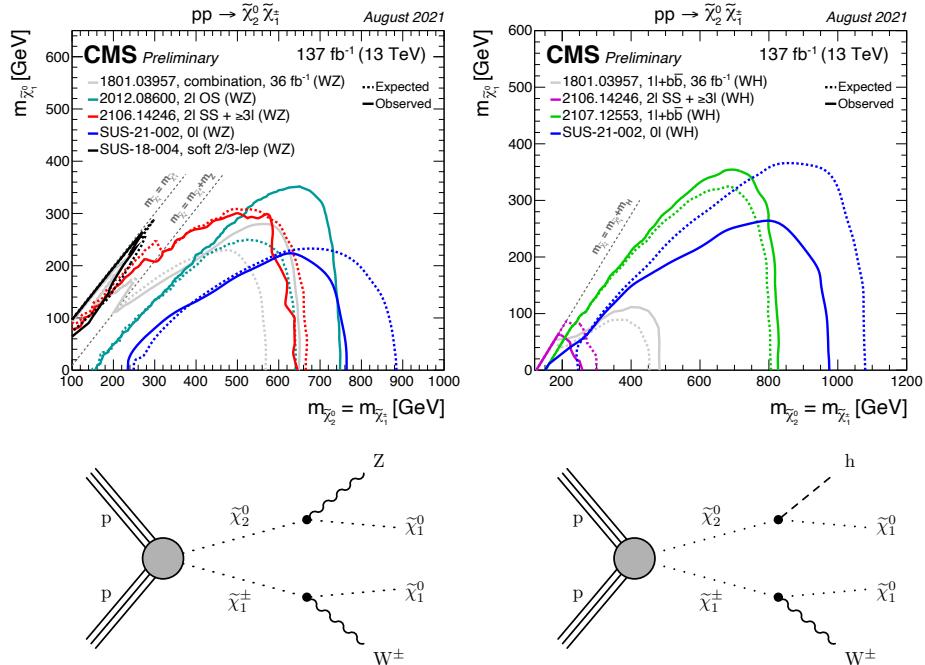


Figure 2.5: Exclusion limits for pair-produced wino-like $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ in SMS models. The corresponding decay topologies are indicated with the diagrams below. Figures from [76], with individual publications indicated in the legends.

2.4 Supersymmetry with Unconventional Signatures

even less stringent due to their lower cross-section, reaching up to 205 GeV for a mass-splitting of 7.5 GeV.

Sleptons – Finally, sleptons have the lowest production cross-section at the LHC. The exclusion limits for pair-produced sleptons are shown in Figure 2.6. The model assumes that the first two slepton families are mass-degenerate, and decays occur as $\tilde{\ell} \rightarrow \ell \tilde{\chi}_1^0$. Considering both left- and right-handed sleptons, masses up to 700 GeV are excluded for light LSPs. Also here, the phase-space with mass-compressed NLSP and LSP is hardly probed and will be discussed more in Chapter 6.

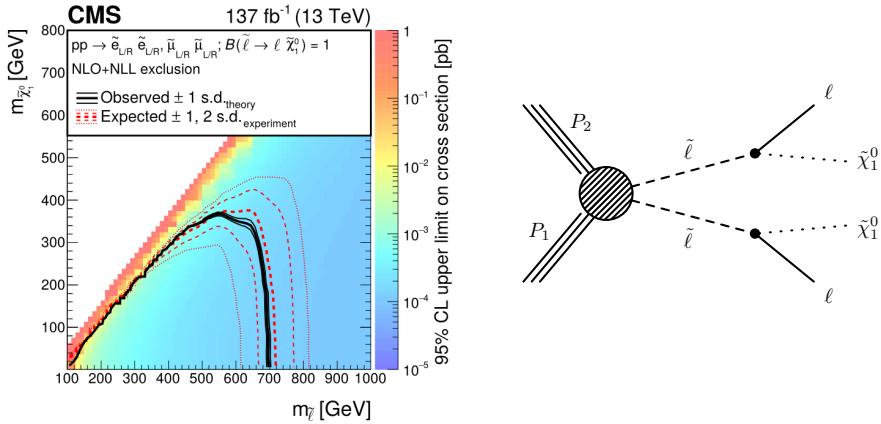


Figure 2.6: Exclusion limits for pair-produced sleptons in SMS models [99]. The sleptons are assumed to be of the first two families, which are generate in mass. The corresponding decay topologies are indicated with the diagram on the right.

2.4 Supersymmetry with Unconventional Signatures

The aforementioned LHC Run 2 results impose tight constraints on realizations of Supersymmetry at the electroweak scale in its simplest form. As such, the general tendency is to assume that the new physics states instead live at higher energy scales, which suggests the need for more powerful collider experiments and higher integrated luminosities. However, there are still several – well motivated – Supersymmetry realizations within reach of the LHC physics potential, that remain largely unprobed due to experimental challenges. Therefore, to give the ultimate answer to the existence of Supersymmetry at the electroweak scale, it is crucial to explore all corners of the phase-space where Supersymmetry may reside. In the following, two Supersymmetry scenarios are highlighted that result in more

Supersymmetry

unconventional signatures, which could have escaped experimental observation so far. Both scenarios play a central role in the search presented in Chapter 5.

2.4.1 Compressed Mass-Spectra

New physics with compressed mass-spectra refers to scenarios in which some of the new particle states are nearly mass-degenerate, their mass-splittings ranging anywhere from $O(10)$ GeV down to sub-GeV level. Within Supersymmetry, there are several scenarios that result in compressed mass-spectra. For example, in case $|\mu| \ll |M_1|, |M_2|$ with μ at the electroweak scale, as motivated by naturalness arguments [100, 101], then $\tilde{\chi}_1^0$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^\pm$ would be nearly mass-degenerate higgsino-like states. Higgsinos are therefore naturally compressed. Alternatively, if $|M_1| < |M_2| \ll |\mu|$, then the above electroweakinos would have a nearly pure wino and/or bino nature. If additionally the $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^\pm$ would be only slightly more massive than the LSP neutralino, which is a thermal relic Dark Matter candidate, then a *coannihilation* process in the early universe could have depleted the LSP population by the amount required to match the observed Dark Matter density of today [102, 103]. Lastly, it should be noted that new physics with compressed mass-spectra can also manifest in other BSM models, outside the framework of Supersymmetry.

The production of mass-compressed electroweakinos at the LHC is particularly hard to target. Due to R-parity conservation, the pair-produced sparticles may each go through a cascade decay into the stable, undetected LSP. The LSP carries away most of the energy due to the small mass-splitting with the (slightly) heavier sparticles. This implies that the remaining SM decay products have only low momentum (referred to as “soft”), typically only a few GeV or less. Furthermore, since the new particles states would be produced nearly at rest, the LSP induces only moderate amounts of p_T^{miss} . Final states therefore consist of only soft leptons or jets and small amounts of p_T^{miss} . This signature is difficult to capture due to many experimental challenges.

Firstly, one challenge is the online event selection. With the overwhelming background of the LHC environment, the absence of high p_T visible decay products and only low p_T^{miss} implies that standard trigger strategies are futile. One way to circumvent this issue is to require a jet from initial state radiation (ISR) in back-to-back configuration with the sparticle pair, giving the latter a boost that induces enough p_T^{miss} for the usage of p_T^{miss} -based triggers. Such event topologies may be further exploited by defining dedicated event variables with the *Recursive Jigsaw*

2.4 Supersymmetry with Unconventional Signatures

Reconstruction [104]. It should be noted however, that the requirement of a hard ISR jet clearly reduces the signal cross-section, which is already low to begin with (see Figure 2.2).

Other challenges arise in the reconstruction and identification of the soft visible decay products. Low p_T jets from sparticle decays are hardly distinguishable from soft QCD background, which motivates the usage of leptonic final states. While the LHC experiments were not designed to target low energy leptons – but instead for high energy signatures from top or Higgs decays – innovative identification strategies both at the trigger level and in the offline selections make final states with soft leptons a promising channel to target mass-compressed electroweakinos. Still, misidentified hadronic activity remains a challenging background that can only be reduced with tighter identification and isolation criteria, which would affect the signal efficiency as well. All these experimental challenges leave the parameter space of Supersymmetry with compressed mass-spectra largely unprobed, and will need to be addressed; In this thesis, the reconstruction and identification of low energy leptons to target mass-compressed new physics signals plays a central role, and will be presented in Chapter 5.

2.4.2 Long-lived Particles

All search results presented in Section 2.3.3 target the production of NLSP pairs that promptly decay to the LSP and SM particles. These searches often require that the final state products originate from the primary vertex interaction via eg. track-matching criteria. However, this leaves them blind to scenarios in which sparticles are long-lived.

If long-lived particles decay within the CMS detector, it may lead to a plethora of unique signatures, consisting of eg. displaced vertices or disappearing tracks, as shown in Figure 2.7. In recent years there has been increased interest to target these type of signatures [105]. Long-lived sparticles are predicted by a range of Supersymmetry models. Examples include long-lived NLSP $\tilde{\chi}_1^0$ or $\tilde{\chi}_1^\pm$ in models of gauge- or anomaly-mediated Supersymmetry breaking, respectively [56, 106]. Furthermore, models of split-Supersymmetry predict long-lived gluinos [57–59, 107]

Finally, a scenario with long-lived *and* mass-compressed Supersymmetric particles is predicted by wino/bino coannihilation in mini-split models [108–110]. This scenario leads to a long-lived wino-like NLSP, which decays to the slightly lighter LSP $\tilde{\chi}_1^0$ – the thermal relic Dark Matter candidate – and soft SM particles. Experimental signatures therefore involve displaced decays and soft visible particles.

Supersymmetry

This is therefore a particularly challenging signature, which is targeted in the search presented in Chapter 5.

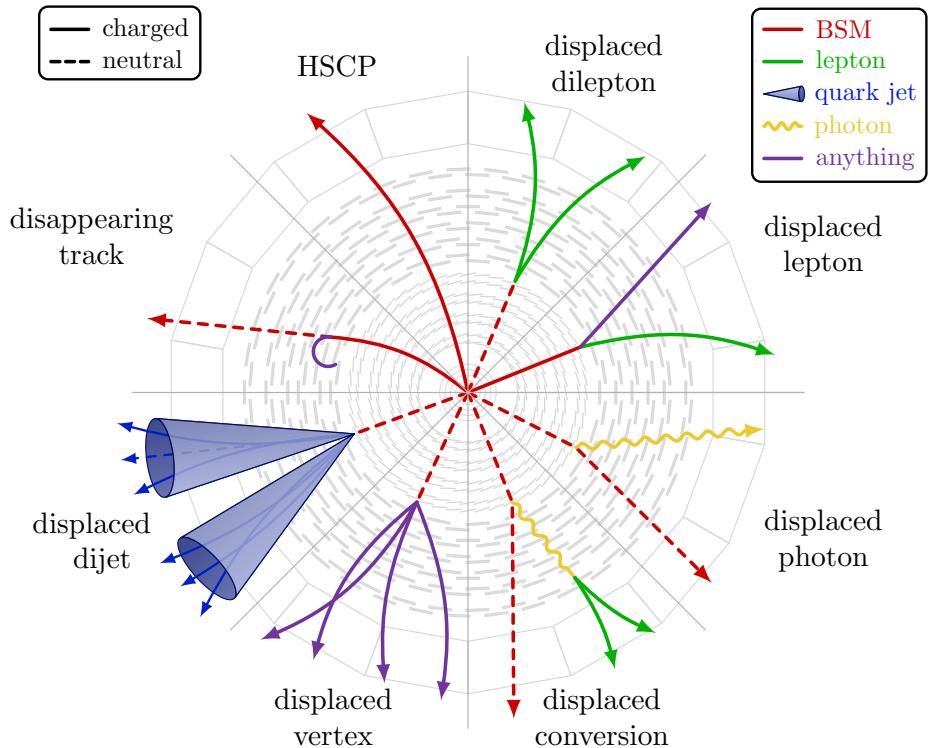


Figure 2.7: Examples of experimental signatures from long-lived particles predicted by BSM theories. HSCP is acronym for *heavy stable charged particle*. Made by Izaak Neutelings.

Chapter 3

The LHC and the CMS Experiment

The European Organization for Nuclear Research (CERN) is one of the world's largest scientific research institutes. Founded in September 1954, it became a unifying force to post-war Europe by serving as a platform for international collaboration towards fundamental physics research. Today, over 600 institutes and universities make use of CERN's facilities, including most notably the world's largest and most powerful particle accelerator: the Large Hadron Collider (LHC). This superconducting circular hadron collider is situated about 100 meters underground, just west of Lake Geneva.

Along its 27 km circumference, the LHC serves a variety of detectors and experiments, the four main ones being: A Toroidal LHC ApparatuS (ATLAS), the Compact Muon Solenoid (CMS), A Large Ion Collider Experiment (ALICE) and the Large Hadron Collider beauty (LHCb) experiment. Together these experiments explore many different aspects of the fundamental nature of matter and the universe. The data analyses and upgrade work presented in this thesis all correspond to the CMS experiment, a general-purpose detector designed to explore the frontiers of energy and luminosity.

This chapter presents a brief description of the LHC and its main characteristics. This is followed by a description of the CMS experiment in Sections 3.2 and 3.3, with an emphasis on the CMS Trigger System. Lastly, Section 3.4 highlights the ongoing developments for the future upgrades of the LHC and the CMS experiment.

3.1 The Large Hadron Collider

The Large Hadron Collider (LHC) [111] follows in the footsteps of the Large Electron-Positron collider (LEP) [112], which operated for 11 years at collision energies up to 209 GeV, and facilitated detailed studies of the electroweak interaction. As of 2000, LEP was dismantled to make way for the LHC, that now operates in the same 27 km long underground tunnels and is designed to run at a center-of-mass energy of 14 TeV and instantaneous luminosity up to $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ for proton-proton collisions. The main motivations for the LHC are summarised as follows.

- **Search for the Higgs boson** – Finding the Higgs boson and studying its properties was the primary target, due to its key role in the mechanism of electroweak symmetry breaking (EWSB). In 2012, both ATLAS and CMS reported the observation of a particle consistent with the SM Higgs, confirming the simplest realization of EWSB in the SM [6, 7].
- **SM precision measurements** – Due to its record-breaking collision energy and luminosity, the LHC is a factory for heavy SM particles, such as W and Z bosons or third generation quarks. The large production rates allow for high precision measurements of particle properties and provide access to rare physics processes.
- **Physics beyond the SM** – Despite its remarkable predictive power, the SM is often viewed as a low energy effective field theory, an approximation of a more fundamental theory, such as Supersymmetry, that comes with new particles and physics processes at higher energy. The high collision energy of the LHC could expose such mechanisms occurring at TeV scale energies.
- **Heavy ion physics** – The LHC can also be used as a heavy ion ($^{208}\text{Pb}^{82+}$) collider, allowing studies of quark-gluon plasma, a plasma of deconfined quarks and gluons also present at the early stages of the universe.

The LHC sits at the end of a large accelerator complex, as depicted in Figure 3.1. To generate proton beams, a source of protons is produced by ionizing hydrogen gas. The protons enter the accelerator chain via the Linear Accelerator (LINAC) after which their energy is increased by the Proton Synchrotron (PS) and Super Proton Synchrotron (SPS) before being injected into the LHC. Through 16 radio frequency (RF) cavities the beams are then accelerated to collision energy, and kept in their circular trajectory by 1232 superconducting dipole magnets. The

3.1 The Large Hadron Collider

LHC beam revolution period of $89 \mu\text{s}$ is divided into 3654 bunches with 25 ns spacing, of which 0(2500) are filled with proton clouds to account for the SPS and LHC injection and beam dump kicker rise times. The particle bunches are made to collide at each of the four large LHC detectors through the use of triplets of quadrupole magnets, at a maximum frequency of 40 MHz.

A key characteristic of the LHC is the center-of-mass energy \sqrt{s} , as this affects the particle production cross-sections (via dependencies in eg. the proton PDF or strong coupling α_s , as discussed in Section 1.4), and defines an upper bound for the mass-scale of emerging particles, including new physics states. During LHC Run 1 (2010-2012), the center-of-mass energy was fixed at 7 and 8 TeV. During Run 2 (2015-2018) this was increased to 13 TeV, and Run 3 (2022-2025) is ongoing at an energy of 13.6 TeV.

Besides the center-of-mass energy, also the instantaneous luminosity L_{inst} is an important property of the LHC machine, since it relates the event rate of a

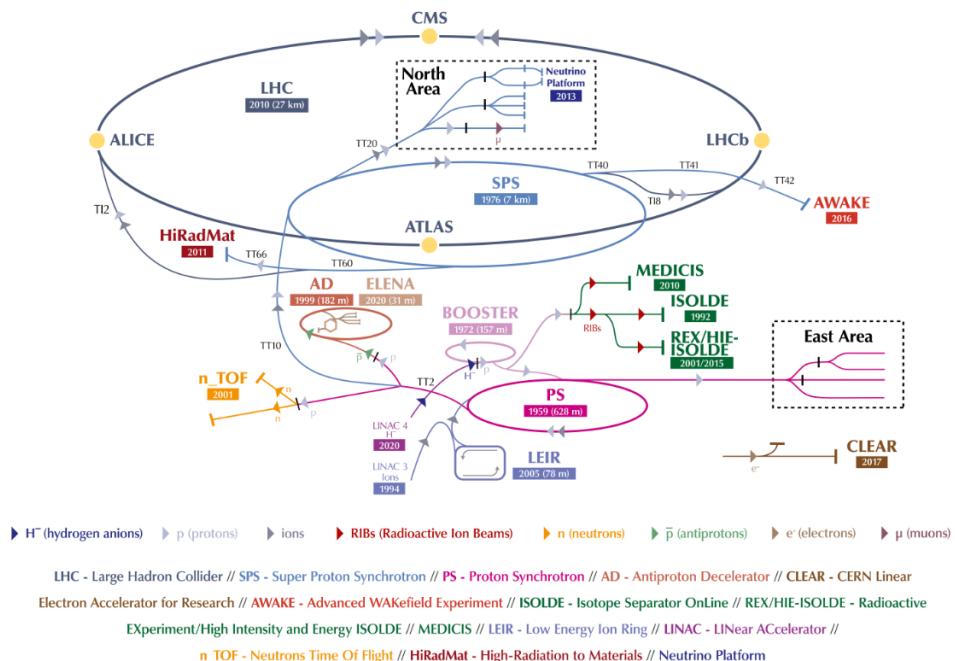


Figure 3.1: Schematic of the CERN accelerator complex [113].

The LHC and the CMS Experiment

process to its respective cross-section σ as:

$$\frac{\partial N}{\partial t} = L_{inst} \times \sigma \quad (3.1)$$

Assuming that the opposite direction LHC beams are identical and have a Gaussian beam distribution, the instantaneous luminosity is characterized by the beam parameters as:

$$L_{inst} = \frac{N_p^2 n_b f_{rev} \gamma_r}{4\pi \epsilon_n \beta^*} R, \quad R = 1/\sqrt{1 + \frac{\theta_c \sigma_z}{2\sigma_t}} \quad (3.2)$$

where N_p is the number of particles per bunch, n_b the number of bunches per beam, f_{rev} the revolution frequency, γ_r is the Lorentz factor, ϵ_n the normalized transverse beam emittance and β^* the β -function characterizing the focus of the beam at the interaction point. R is a geometric luminosity reduction factor that accounts for the fact that the particle bunches with longitudinal and transverse RMS sizes σ_z and σ_t , respectively, cross under a non-zero angle θ_c . The LHC design values are shown in Table 3.1 (Section 3.4.1).

The total amount of collision data accumulated over a certain span of time is called the integrated luminosity L , typically expressed in units of inverse femtobarn ($1 \text{ fb}^{-1} = 10^{-39} \text{ cm}^{-2}$). The integrated luminosities delivered by the LHC and collected by the CMS experiment are shown in Figure 3.2 (top).

Finally, as the LHC beams are circulating and crossing the interaction point, multiple protons per bunch-crossing are made to collide. The vast majority of collisions are *soft*, with only small amounts of momentum transfer between protons at relatively large distance. These interactions are typically not energetic enough to create massive particles and the proton (remnants) escape down the beam pipe. Only a small fraction of the collisions are *hard* interactions, with enough momentum transfer between the proton constituents to create massive particles. The number of hard interactions per bunch-crossing (also known as *pileup*) recorded by CMS is shown in Figure 3.2 (bottom).

3.2 The CMS Detector

The CMS detector [116] is a general-purpose detector operating at the LHC, designed to explore the energy and luminosity frontiers. It seeks to fulfill the aforementioned physics goals by collecting high quality data, which translates into an extraordinarily challenging set of requirements on the detector subsystems. This

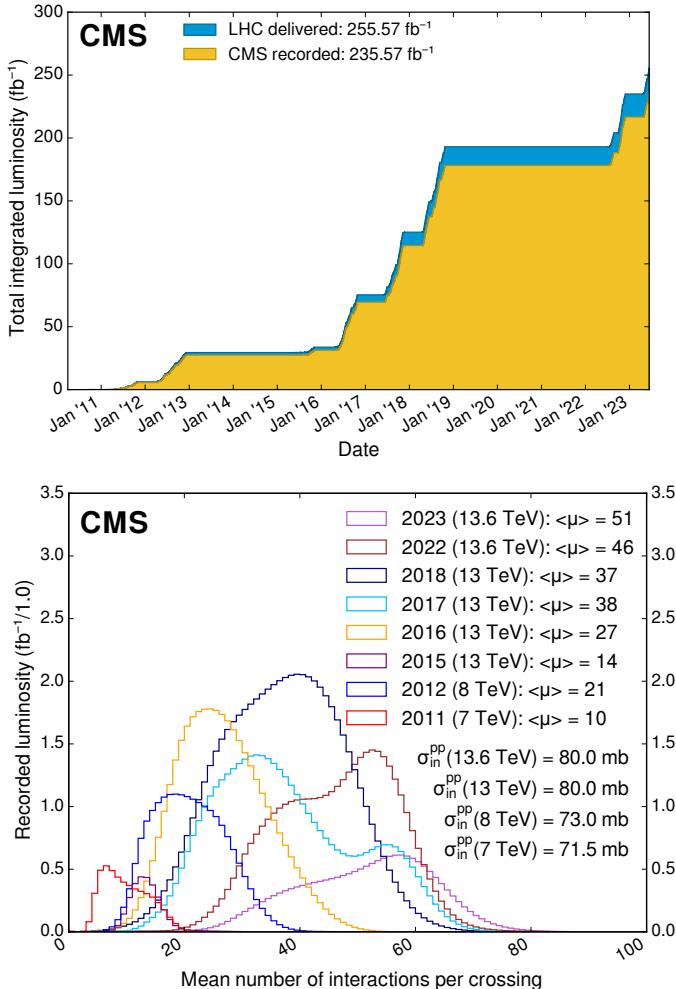


Figure 3.2: Top: Cumulative luminosity from proton-proton collisions delivered by the LHC and recorded by CMS. Bottom: Mean number of pileup interactions per bunch-crossing in CMS [114].

The LHC and the CMS Experiment

includes fast and efficient online event selection (triggering), excellent momentum, energy and mass resolution, particle identification and radiation hardness. The detector is built in an onion-like fashion, with cylindrical layers coaxial to the beam axis and perpendicular end-cap disks to maximally enclose the interaction point located at the center. A schematic of the detector is shown in Figure 3.3, including its most characterizing feature: the high-field superconducting solenoid magnet.

The origin of the coordinate system of CMS is fixed at the nominal interaction point, with the x-axis pointing towards the center of the LHC, the y-axis pointing upwards and the z-axis aligned with the beam-axis and pointing to the Jura mountains. Alternatively, a cylindrical coordinate system is often used instead, conform the shape of the detector, where $r = \sqrt{x^2 + y^2}$ is the radial distance from the interaction point, ϕ is the azimuthal angle from the x-axis in the x-y plane, and θ is the polar angle from the z-axis in the y-z plane. The coordinate system is illustrated in Figure 3.4. In particle collisions the z-component of the momenta p_z of the interacting partons is unknown. However, the transverse component $p_T = \sqrt{p_x^2 + p_y^2}$ is negligible, which means that also the sum of p_T of the outgoing

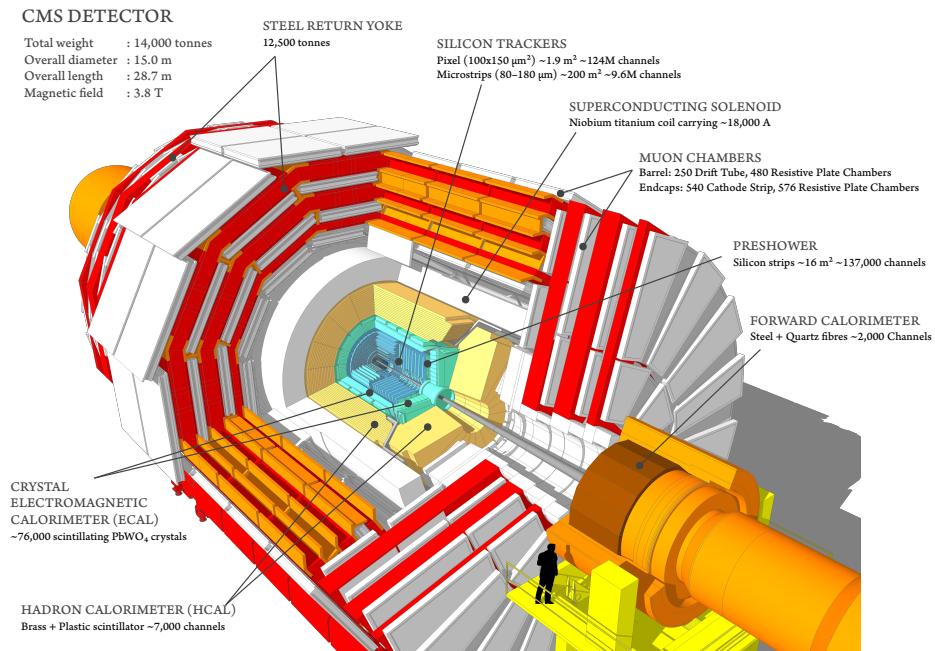


Figure 3.3: Layout of the CMS detector [115].

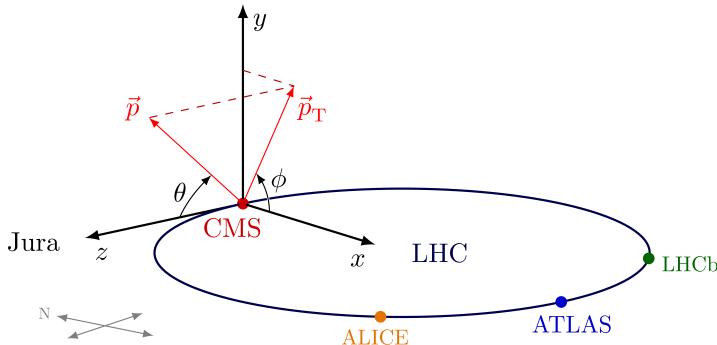


Figure 3.4: Illustration of the CMS coordinate system [117].

particles from the interaction must be zero. This can be used to infer the momenta of particles that do not interact with the detector (such as neutrinos or new physics states). Lastly, to remain invariant under Lorentz boosts along the z -axis, the pseudorapidity $\eta = -\ln(\tan(\theta/2))$ may be used, which is an approximation of the rapidity $y = \frac{1}{2} \left(\frac{E+p_z}{E-p_z} \right)$ for ultra-relativistic particles with mass $m \ll E$.

3.2.1 The Solenoid Magnet

The superconducting NiTb solenoid magnet [118] is a central feature of the CMS detector. It has a free bore diameter of 6 m and length of 12.5 m, and encloses the Inner Tracker and the Calorimeters. With the solenoid itself being housed inside a vacuum tank to maintain the operating temperature of 4 K, it can reach a central magnetic flux density of up to 4 T inside the coil, as shown in Figure 3.5, storing 2.6 GJ at full current. A magnetic field this powerful brings substantial improvements to the detector performance, in particular to the muon tracking, as it increases momentum resolution and results in higher performance muon triggers because of a sharper trigger turn-on. Furthermore, it allows for a more compact tracker system. Outside the solenoid, a 12000-tonne iron yolk structure is used to increase the field homogeneity and return the magnetic flux of the solenoid. The yolk consists of three barrel wheels and three endcap disk at each end, interleaved with the muon chambers. The thickness of the return yolk contributes to its absorbing capabilities, allowing only weakly interacting particles to pass, which in turn allows for safe muon identification. Furthermore, because of its excessive weight, its size and rigidity, the magnet-yolk structure serves as the principle support for the other subsystems of CMS.

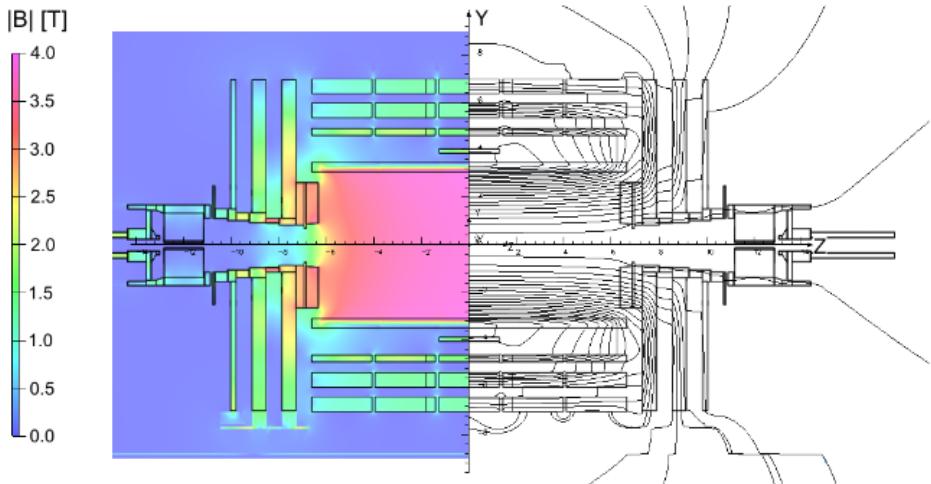


Figure 3.5: Map of the magnetic field $|B|$ (left) and field lines (right) in the CMS detector [119].

3.2.2 The Inner Tracker

The innermost sub-detector of CMS is the Inner Tracker [120], with its dimensions extending up to $r < 120$ cm and $|z| < 270$ cm. One of its main purposes is to efficiently and precisely measure the trajectories of charged particles and reconstruct interaction vertices. In particular the reconstruction of isolated leptons was a key quality of the Inner Tracker design, as it would allow for suppression of $t\bar{t}$ and $Zb\bar{b}$ backgrounds to a level needed for observation of $H \rightarrow ZZ^{(*)} \rightarrow 4l^\pm$. Furthermore, lepton isolation criteria play an important role in many searches for new physics, for example those exploring the electroweak sector of Supersymmetry, as will be discussed in Chapters 5 and 6. On the other hand, accurate reconstruction of (secondary) interaction vertices is a capability that is essential to the reconstruction and identification of b jets and τ decays, signatures that are crucial to studies involving top physics, CP violation, new physics searches and more.

The Inner Tracker was designed to perform these tasks in the high luminosity environment of the LHC, implying thousands of particles traversing the tracker volume at each bunch-crossing (ie. every 25 ns). The large particle flux demands high granularity sensors to maintain low cell occupancy in addition to fast charge collection. The Inner Tracker is therefore entirely based on silicon detector technology, providing sufficient granularity, readout speed and radiation hardness. The full Inner Tracker is operated at a temperature of -10 °C to cope with the

high power density, while satisfying tight constraints on the material budget to minimize bremsstrahlung, photon conversions and hadronic interactions.

The Inner Tracker consists of two multilayered sub-detectors: a pixel detector located closest to the interaction region, and a strip detector, as is shown in Figure 3.6. The pixel detector was upgraded during the end-of-year technical stop of LHC in 2016/2017 [122, 123] and now consists of 4 concentric, cylindrical layers in the barrel region (BPIX) and 3 disks in the end-cap region (FPIX), aiming to maximize hit redundancy over the range $|\eta| < 2.5$. The layers in BPIX have a length of 548.8 mm, are located at $r=30, 68, 109$ and 160 mm, and contain a total of 79 M pixels, each sized $100 \times 150 \mu\text{m}^2$. The FPIX layers, carrying another 46 M pixels, provide radial coverage between $r=45$ and 161 mm, and are located at $z=191, 396$ and 516 mm. They are slightly rotated to achieve optimal resolution in the azimuthal and radial directions.

The pixel detector is enclosed by the strip detector, which has lower constraints on granularity and radiation hardness due to smaller particle flux compared to the pixel detector. The Tracker Inner Barrel (TIB) extends up to $r=550$ mm and consists of 4 layers, while the Tracker Inner Disk (TID) has 3 disks in each end-cap. The Tracker Outer Barrel (TOB) extends up to $r=1160$ mm and consists of 6 layers, while the Tracker Endcaps (TEC) are made of 9 disks each, extending up to $|z|=2820$ mm. Up to 4, 6 and 9 $r\phi$ measurements are delivered by the TIB/TID, TOB and TEC systems, respectively. The total strip detectors contains 9.3 M strips of varying size and geometry, together yielding 198 m^2 of active silicon area.

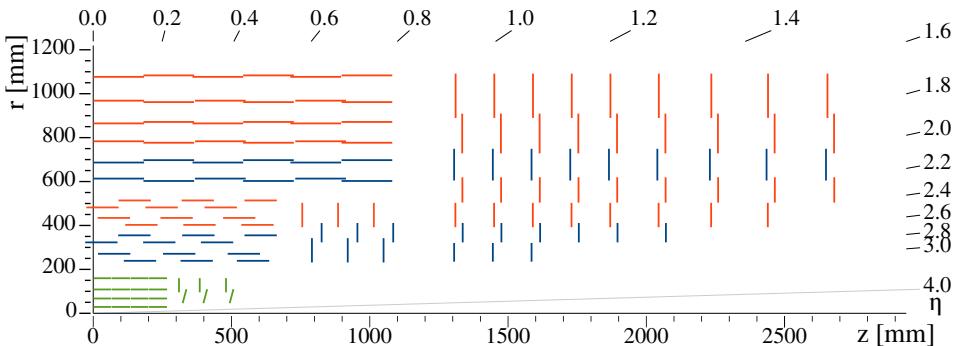


Figure 3.6: Layout of one quarter of the Phase-1 CMS Tracker in the r - z plane (after the pixel upgrade of 2016/2017) [121]. The pixel detector is shown in green, while the red and blue segments correspond to single-sided and double-sided strip modules.

3.2.3 The Electromagnetic Calorimeter

The design of the Electromagnetic Calorimeter (ECAL) [124] was mainly driven by the ability to observe $H \rightarrow \gamma\gamma$, putting tight requirements on the experimental di-photon mass-resolution, which entirely governs the width of the Higgs signal. To obtain sufficient granularity with a system that is small enough to fit inside the solenoid magnet, the ECAL is made of lead tungstate ($PbWO_4$) crystals. The high density (8.28 g/cm^3) and short radiation length ($X_0=0.89 \text{ cm}$) of $PbWO_4$ make for a compact calorimeter, while its small Molière radius (2.19 cm), ie. the lateral spread containing 90% of the energy deposit of an electromagnetic shower, allows for sufficient granularity. Furthermore, the short scintillation decay time of the $PbWO_4$ crystals implies that roughly 80% of the scintillation light is collected every 25 ns, making them fast enough to match the LHC bunch-crossing frequency of 40 MHz.

The ECAL is divided into a barrel region (EB) that extends up to $|\eta| < 1.479$, and two end-cap regions (EE) covering $1.479 < |\eta| < 3.0$. A schematic is shown in Figure 3.7. The EB consists of 61200 trapezoidal shaped crystals, each with a length of 23 cm ($25.8 X_0$) and covering $\Delta\eta \times \Delta\phi = 0.0174 \times 0.0174$, which corresponds to roughly $22 \times 22 \text{ mm}^2$ at the front (matching the Molière radius) and $26 \times 26 \text{ mm}^2$ at the rear. Each EE consists of 7324 crystals that are slightly broader ($28.6 \times 28.6 \text{ mm}^2$ in the front, $30 \times 30 \text{ mm}^2$ in the rear) and shorter (220 mm or $24.7 X_0$) than in the barrel region. The scintillating light of crystals in the EB

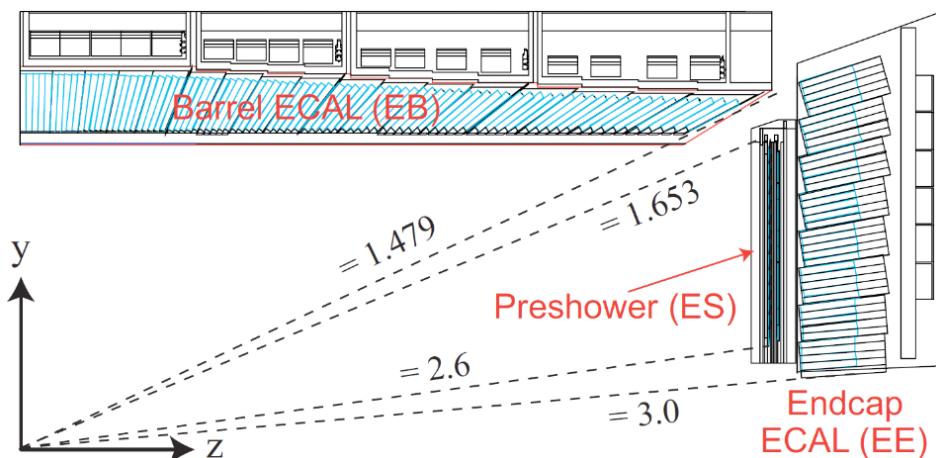


Figure 3.7: Layout of one quarter of the CMS Electromagnetic Calorimeter [125].

and EE are collected by avalanche photodiodes (APD) and vacuum phototriodes (VPT), respectively.

In each end-cap an additional preshower (ES) detector is placed, covering $1.653 < |\eta| < 2.6$. It consists of lead radiators to initiate electromagnetic showers from incoming photons and electrons, and two orthogonal layers of silicon strip sensors to measure the energy deposits and transverse shower shapes. The ES mainly serves to distinguish neutral pions and photons, or charged pions and electrons, but also contributes to improved position measurements due to its high granularity. It is therefore complementary to the EE.

The relative energy resolution, defined as the ratio of the standard deviation σ_E of a Voigtian profile, and the mean energy E , is typically measured with $Z \rightarrow ee$ decays and parameterized as:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2 \quad (3.3)$$

where the stochastic term S accounts for event-to-event fluctuations in the lateral shower containment, the number of photoelectrons and energy deposits. The noise term N includes contributions from electronics, digitization and pileup, and the constant term C accounts for crystal non-uniformity and calorimeter calibrations. The total energy resolution at $\eta=0$ and 2 was measured to be $(\sigma_E/E)=2\%$ and 4%, respectively, which is comparable to that of Run 1, despite the increased LHC luminosity and ageing effects of the detector [126].

3.2.4 The Hadronic Calorimeter

The Hadronic Calorimeter (HCAL) [127] plays a crucial role in the reconstruction of neutral hadrons and missing transverse energy induced by eg. neutrinos or WIMPs. It consists of four sub-systems: the barrel (HB), endcap (HE), outer (HO) and forward (HF) calorimeters, as shown in Figure 3.8.

The HB and HE are sampling calorimeters located between the ECAL and solenoid magnet at $1.77 < r < 2.95$ m, and consist of alternating layers of brass absorbers and active scintillator material. The HB covers $0 < |\eta| < 1.4$, while HE covers $1.3 < |\eta| < 3.0$. Due to the orientation of the transition region between HE and HB ($1.3 < |\eta| < 1.4$), no projective gap occurs, which maximizes the hermicity needed for measurements of the missing transverse energy. The lateral granularity is chosen to be $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ at $|\eta| < 1.6$ and $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$ at $|\eta| > 1.6$ to conveniently match the crystal geometry of the ECAL.

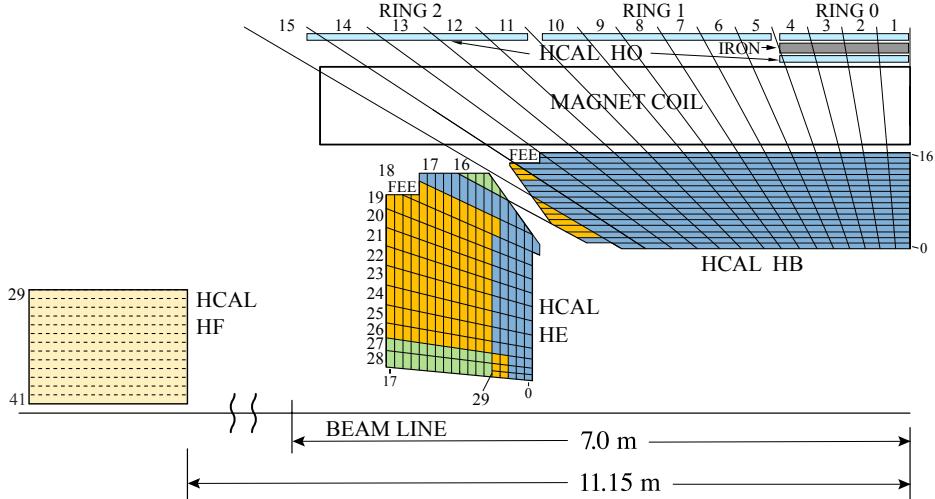


Figure 3.8: Schematic of one quarter of the CMS HCAL detector [128].

Due to space constraints, the thickness of the HB is limited to $5.39\lambda_I$ (nuclear interaction length $\lambda_I=16.42$ cm) at $\eta=0$. This increases as $1/\sin(\theta)$ to $10.6\lambda_I$ at $|\eta|=1.3$. Since the amount of stopping power in the central region does not sufficiently contain hadron showers, the HB is complemented with the HO that is constructed outside the solenoid magnet to “catch the tails” of the shower energy. It consists of several scintillator layers and together with the solenoid as additional absorber extends the depth of the HB to a minimum of $11.8\lambda_I$.

The forward regions of the detector suffer from by far the highest particle fluxes, with roughly 1 Grad ($=10^7$ Gy) radiation expected at $|\eta| = 5$ after an integrated luminosity of 500 fb^{-1} . To maximize hermicity while withstanding this hostile environment, two HF calorimeters, weighing 250 ton each, are placed at 11.15 m from the interaction point and cover $3.0 < |\eta| < 5.0$. They are made of steel absorbers with quartz fibers as active medium, allowing measurements of hadronic jets with energies up to several TeV through the detection of Cherenkov light.

3.2.5 The Muon system

Muons are less affected by radiative losses than electrons, resulting in unique signatures that are relatively easy to detect and allow for high resolution measurements. Clean muonic final states played an important role in the observation of

$H \rightarrow ZZ^{(*)} \rightarrow 4l$ and are expected to have high discovery potential for Supersymmetry or other new physics realizations. The CMS muon spectrometer [129], its importance emphasised by the name of the experiment, serves three purposes: muon identification, muon momentum measurements (both stand-alone and together with the Inner Tracker), and triggering on muons.

The muon system consists of 3 sub-systems, each employing a different type of gaseous detector technology: The Drift Tube (DT) chambers, Cathode Strip Chambers (CSC) and Resistive Plate Chambers (RPC), together covering $|\eta| < 2.4$ with no acceptance loss. A schematic of the muon spectrometer is shown in Figure 3.9. A fourth sub-detector, part of the Gas Electron Multiplier (GEM) project [131], was installed in 2019 in preparation of data-taking at even higher luminosity. This system will be discussed in Section 3.4.2.

The 250 DT chambers are divided over 4 cylindrical layers, interleaved with the magnet return yoke and cover $|\eta| < 1.2$. They contain a total of 172 k rectangular drift cells that are made of a cathode housing filled with a gas mixture (85% Ar, 15% CO₂) that ionizes with incoming muons. The generated electrons drift to the anode wire at the center of the cell, resulting in an electrical signal, and the drift

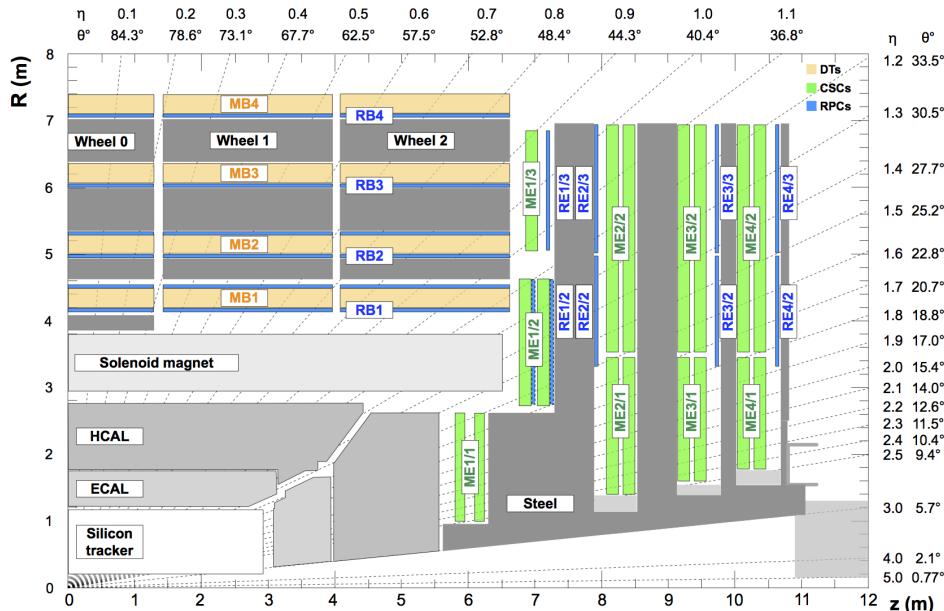


Figure 3.9: Layout of one quarter of the muon spectrometer [130].

The LHC and the CMS Experiment

time is used to infer of the muon position. Within each DT chamber, multiple layers of drift tubes are stacked to resolve any left-right ambiguity and measure the muon coordinates in both the (r, ϕ) and (r, z) planes.

In the end-cap regions the larger occupancy and stronger, non-uniform magnetic field demand a different detection technology. CSCs are employed to cover $0.9 < |\eta| < 2.4$ due to their more suitable segmentation, radiation tolerance and magnetic field compatibility. Each CSC contains 6 detection layers, each consisting of two parallel cathode planes filled with a gas mixture (30% Ar, 50% CO₂, 10% CF₄). One cathode per layer is divided into strips along the r -direction (with constant $\Delta\phi$), running orthogonal to the anode wires at the center of each layer. The charge distributions at the finely segmented strips and anode wires, included by incoming muons, allows for an accurate measurement of the muon position in the $r - \phi$ plane.

The DTs and CSCs are complemented by the RPCs in the range $|\eta| < 1.9$, which are mainly designed for trigger purposes. In particular, due to their timing resolution of less than 3 ns, the RPCs are able to unambiguously associate muons to their respective LHC bunch crossings. The RPCs are constructed from two highly resistive, parallel Bakelite plates, enclosing yet another gas mixture (96% C₂H₂F₄, 3.5% i-C₄H₁₀, 0.5% SF₆) that is optimized for fast ionization and charge collection. Aluminum strips outside the bakelite chambers are used for readout and provide also a coarse measurement of the muon position.

3.3 The CMS Trigger System

At a maximum bunch-crossing rate of 40 MHz, and with each collision event resulting in roughly 1 MB of raw CMS data, tremendous data rates are expected. The implications for detector readout, data-processing and storage are far beyond technical feasibility. Therefore, CMS uses a two-tiered trigger system to reduce the event rate down to 1 kHz, selecting the events to be written to permanent storage based on their physics contents:

- The Level-1 Trigger (L1T) [132, 133] is based on custom hardware processors and serves to reduce the event rate from 40 MHz to 100 kHz using a simplified event reconstruction.
- The High-Level Trigger (HLT) [134] consists of a CPU farm and further reduces the rate down to 1 kHz based on a more detailed event description. The events selected by the HLT are transferred to the Tier 0 data-center [135]

for *prompt reconstruction* within 48 hours from data-taking and permanently stored.

Below follows a description of the L1T and HLT as operated for the standard data-streams, which serve the core CMS physics program. Additionally, the trigger and data acquisition system is capable of *data scouting* and *data parking*, which aim to circumvent the trigger bandwidth limitations in order to maximize the amount of physics data available for analysis [136]. In data scouting, only trigger-level information is stored instead of the full event content, which decreases the event size, hence allowing an increased output rate. In data parking, the data are saved without performing the prompt reconstruction (which is what constrains the HLT output rate). This way, additional rates of raw data can be stored and reconstructed when computing resources are available.

3.3.1 The Level-1 Trigger

At each bunch-crossing, ie. every 25 ns, a high-resolution snapshot of the full CMS detector is added to the pipelined memories in the on-detector electronics, which has a capacity of 160 bunch crossings. Therefore, from the moment a given collision event occurs, there is only a futile 4 μ s to determine whether to reject the event or accept it for further evaluation by the HLT. This decision making is

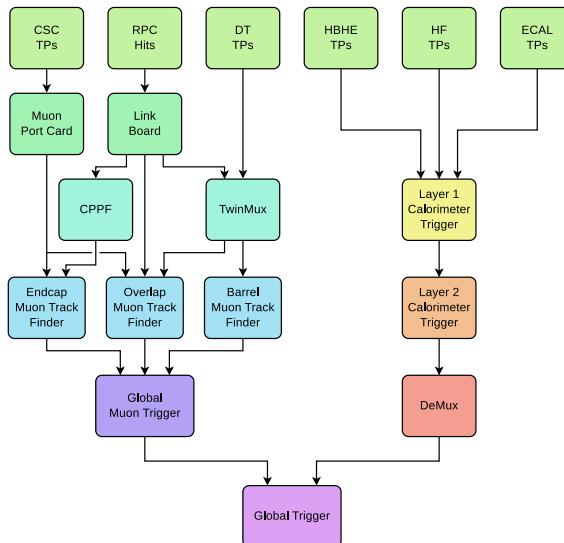


Figure 3.10: Schematic of the CMS L1T architecture during Run 2 [137].

performed by the Level-1 Trigger that consists of custom-designed hardware, with electronics located both on-detector and in the service cavern, approximately 90 m away from the detector. The maximum latency of $4\ \mu\text{s}$ limits the amount and granularity of the data that can be used by the L1T. In practice this implies that the accept-reject decision has to be made based on coarsely segmented data from the Calorimeter and the Muon spectrometer only. The L1T architecture, shown in Figure 3.10, is therefore divided into a Calorimeter Trigger and a Muon trigger, which construct low resolution physics objects (candidates): L1 e/γ , taus, jets and energy sums and muons. These objects are then sent to the Global Trigger, which makes the L1 trigger decision.

The Calorimeter Trigger

The data-flow of the Calorimeter Trigger [138, 139] starts with the generation of trigger primitives (TPs) in the on-detector electronics boards. These represent the energy deposits in the calorimeters, which in the barrel region are segmented into trigger towers (TT) of 5×5 EB crystals grouped with the HB tower directly behind, corresponding to $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$. In the end-caps the EE crystals and deposits in HE and HF are grouped into TTs with sizes up to $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$. The TPs are sent to the Layer 1 Calorimeter Trigger (CALO-L1), which consists of 18 Calorimeter Trigger Processor 7 (CTP7) cards, each spanning 4 out of 72 TTs in the ϕ direction and the full η range. The main task of the CALO-L1 cards is to calibrate the TPs, in order to compensate for particle energy loss in the tracker material and account for changing calorimeter response due to eg. radiation damage. The calibrated TP information, consisting of the summed energy deposit in the ECAL and HCAL, the ECAL/HCAL energy ratio and quality flags, is then sent to the Layer 2 Calorimeter Trigger (CALO-L2).

The CALO-L2 reconstructs and further calibrates calorimeter-based physics objects (e/γ , taus, jets and energy sums). It consists of 9 Master Processor 7 (MP7) cards, each receiving the information from all CALO-L1 cards (ie. spanning all TTs), in order to run with a 9-fold time multiplexing. Each MP7 card therefore has access to the full L1 calorimeter information of every 9th event, which removes regional boundaries in object reconstruction and facilitates more sophisticated algorithms based on a global view of the calorimeter. The e/γ reconstruction is seeded by a TT that constitutes a local energy maximum with $E_T > 2\ \text{GeV}$. Clusters with up to 8 neighbouring TTs are built dynamically to contain the majority of the particle energy, while limiting the contribution from pileup. The cluster extends in the ϕ direction to account for bremsstrahlung effects

and a 6×9 TT region around the seed TT is used to define the e/γ isolation. A schematic of the TT geometry is shown in Figure 3.12 (left). Electrons and photons cannot be distinguished, due to the absence of tracking information in the Phase-1 L1T, but discrimination against hadron-induced showers can be achieved via the ECAL/HCAL energy ratio and the compactness of the cluster. The reconstruction of hadronically decaying taus (τ_h) relies on the same clusters, but allows merging of closeby clusters, reflecting the one, two or three prong τ_h -decays to charged or neutral pions. The isolation and number of clustered TTs are important quantities to distinguish τ_h objects from QCD-induced jets, that generally have more surrounding hadronic activity. Finally, the L1 jet reconstruction is based on a 9×9 TT region centered around the seed TT that constitutes a local energy maximum with $E_T > 4$ GeV, as shown in Figure 3.12 (right). In the barrel region this matches the 0.4 clustering size used by the offline anti- k_T jet reconstruction [140]. A *chunky donut*, four 3×9 TT regions around the jet, is defined to estimate contributions from pileup, and the jets are further calibrated to ensure

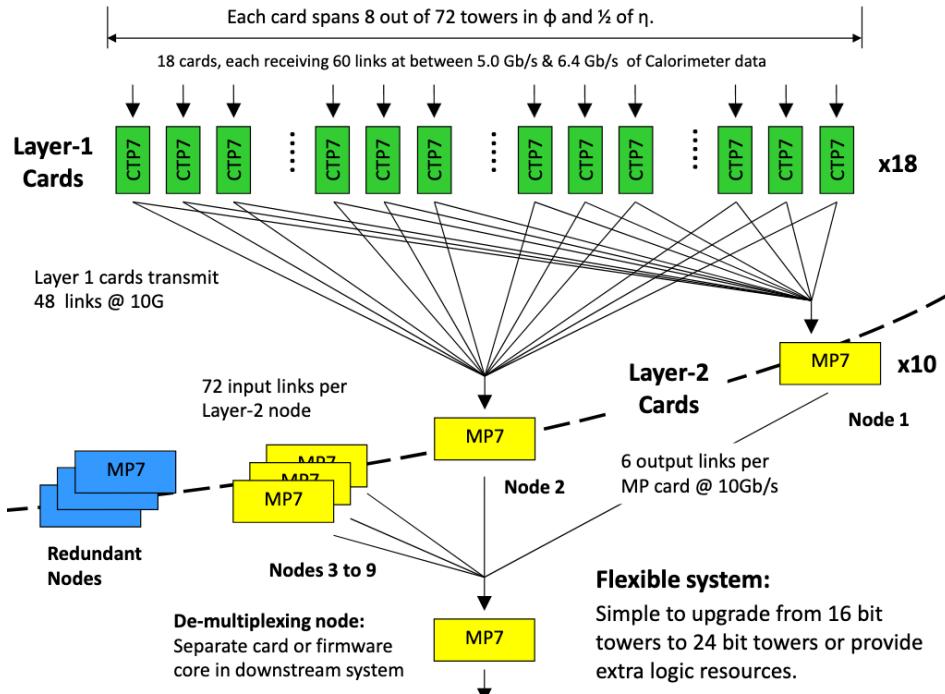


Figure 3.11: Schematic of the Calorimeter Trigger architecture [139].

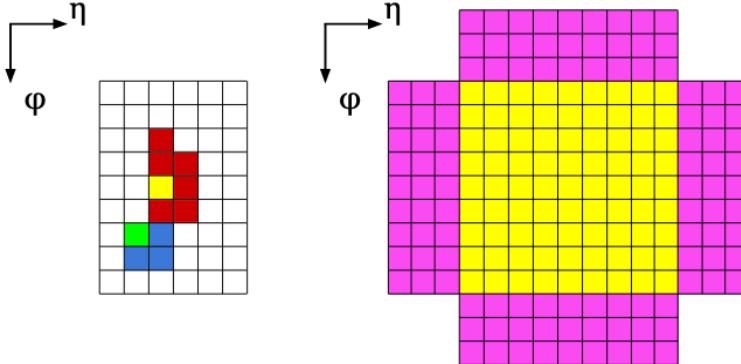


Figure 3.12: Left: Example clustering of trigger towers for the L1T e/γ and τ_h reconstruction, with each square representing one tower. The yellow and green towers are seeds for the red and blue clusters, respectively. The τ_h algorithm allows merging of nearby clusters, and the 6×9 area around the (yellow) seed of the main cluster marks the isolation region. Right: 9×9 tower region used for L1T jet reconstruction, with a surrounding *chunky donut* to estimate pile-up contributions.

consistent jet energy response. The collection of all jets in the event is used to compute the scalar sum of transverse jet energies (H_T) and the calorimeter-based missing transverse energy (E_T^{miss}).

After the object reconstruction by the time-multiplexed CALO-L2, the events are collected by the demultiplexer (DeMux) board, which serialized and formats the data before sending them to the Global Trigger.

The Muon Trigger

The L1 Muon Trigger [141, 142] relies on a regional based approach. Three track finders are used to reconstruct muons, each covering a distinct η -range and combining information from the relevant muon subdetectors (DT, RPC, CSC). The trigger primitives from the CSC and DT systems are formed with on-detector electronics and consist of track segments, which are the result of a fit through the hits in the various layers within one CSC or DT chamber. The TPs from the RPC system are simply the RPC hits themselves. All TPs carry coordinate, timing and quality information.

The Barrel Muon Track Finder (BMTF) covers the region $0 < |\eta| < 0.83$ and receives inputs from the DT and RPC subdetectors. However, before sending the TPs to the BMTF, they are first combined into superprimitives (SPs) by a TwinMux board to exploit the high position and timing resolution of the DT and

RPC systems, respectively. Independently in 12 ϕ -wedges of 30° each, the BMTF then forms pairs of compatible SPs by extrapolating the trajectory from an inner station to outer station based on the SP bending angle and quality. Pairs of SPs from all barrel layers are then combined to form the L1 barrel muon candidates and the p_T , η and ϕ coordinates are assigned using look-up tables (LUTs).

The region $0 < |\eta| < 1.23$ is covered by the Overlap Muon Track Finder (OMTF), which receives unmerged DT and RPC TPs from the TwinMux as well as TPs from the CSCs. The high redundancy of muon stations in this region yields 18 layers that can be used for L1 muon reconstruction. The reconstruction is performed independently in 6 ϕ -sectors of 60° , on both sides of the detector, and is seeded by a single reference TP in one of the layers, prioritizing TPs from one of the inner layers and with high ϕ resolution. Due to the more complex detector geometry and non-uniformity of the magnetic field, compatible hits in other layers are found via 52 predefined *golden patterns*, corresponding to 2 charge- and 26 p_T -hypotheses in the range of $2 < p_T < 140$ GeV.

Finally, the Endcap Muon Track Finder (EMTF) covers $1.23 < |\eta| < 2.4$, receives inputs from the RPCs and CSCs, and - similar to the OMTF - is segmented into 6 ϕ -sectors of 60° per endcap, with muon reconstruction relying on predefined patterns. The p_T assignment in the forward region is more complicated due to interactions with the detector material and irregularities in magnetic field. A boosted decision tree (BDT) regression technique, trained on simulated single-muon events, is therefore employed to estimate the muon p_T . The BDT output is pre-evaluated for all combinations of input values and stored in LUTs to allow for fast p_T determination.

The BMTF, OMTF and EMTF each send up to 36 muon candidates to the Global Muon Trigger (μ GMT). The μ GMT resolves any duplicates resulting from the $20\text{--}30^\circ$ extra margin per ϕ -sector that is assigned at reconstruction to account for muon bending. Furthermore, the μ GMT sorts the candidates based on muon p_T and quality and finally sends the 8 best candidates to the Global Trigger.

The Global Trigger

All physics object that the Global Trigger (μ GT) received from the Calorimeter and Muon Triggers are compared against a *menu* of trigger seeds. Each seed defines an event topology consisting of one of more L1 objects, and the menu is chosen to facilitate a broad and diverse physics program. The result is propagated to the Timing and Control Distribution System (TCDS), which - as the name implies - is responsible for the distribution of the timing signals (such as the LHC

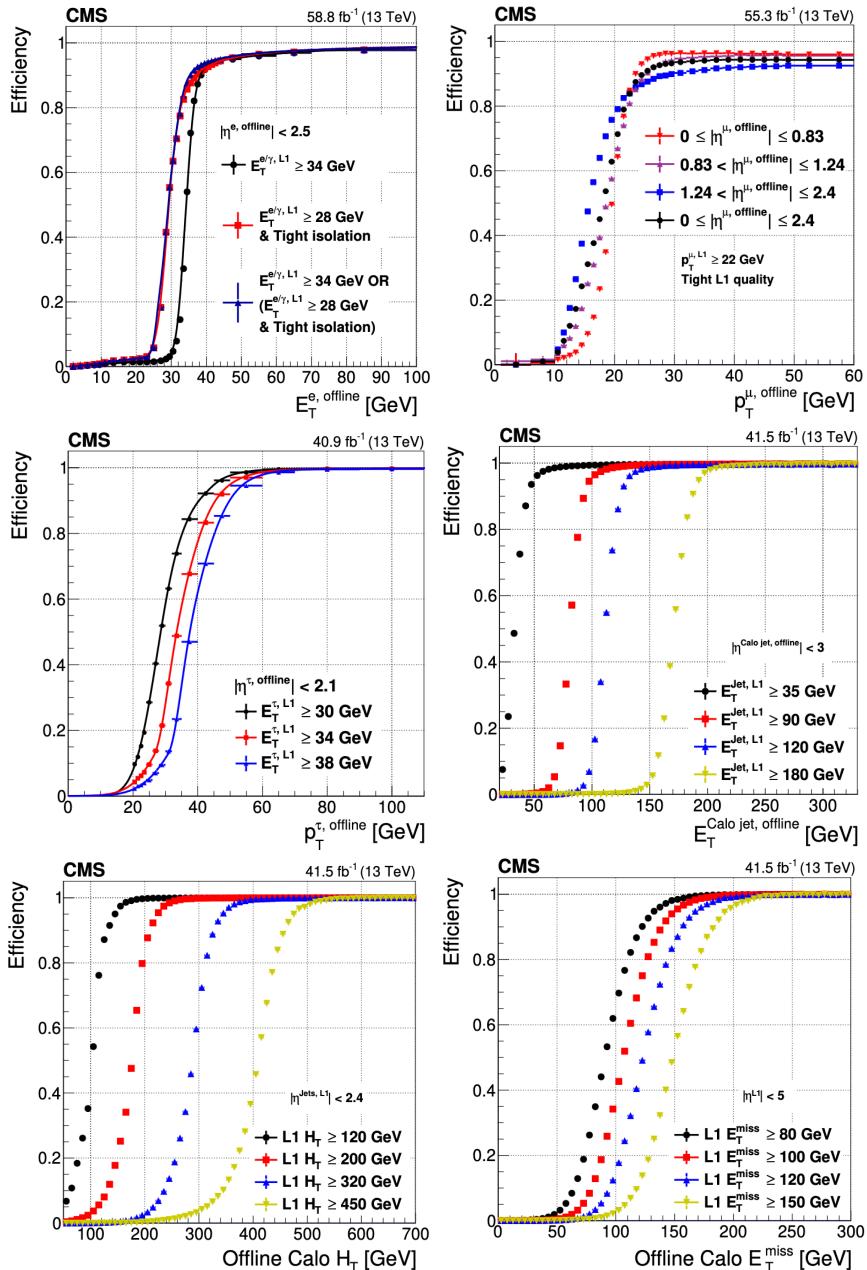


Figure 3.13: Run 2 L1T efficiency curves for representative trigger seeds [137].

clock) and control commands throughout the detector. It keeps track of trigger suppression factors that may lead to operational deadtime, such as non-readiness of front-end detectors (FEDs), bunch masks, data acquisition backpressure, etc., and can throttle the trigger to avoid loss of synchronization. The TCDS itself also generates triggers that serve monitoring and calibration purposes. Once an event is selected by any of the physics, random or calibration seeds, and it is not suppressed by the aforementioned reasons, a Level-1 Accept signal is sent from the TCDS back to the FEDs, initiating full detector readout and propagation of the event data to the HLT.

The performance of the L1T during Run 2 is evaluated in Ref. [137]. and the efficiencies with respect to offline reconstructed physics objects for the most used L1T seeds are shown in Figure 3.13.

3.3.2 The High Level Trigger

The High Level Trigger (HLT) receives event data at a maximum rate of 100 kHz and is tasked to reduce this further down to 1 kHz. In contrast to the L1T, the HLT has access to the entire event information - including data from the Tracker - and is run on a farm of $O(30k)$ CPU cores. This allows the HLT to run the full event reconstruction as it is performed offline, modulo a few minor changes to stay within the maximum allowed latency of approximately 300 ms, such as a simplification of the CPU-intensive track reconstruction. The inclusion of tracking information and the offline-like event reconstruction at HLT based on more granular data gives access to higher level objects, including track-matched muon and calorimeter objects (allowing eg. discrimination of electrons against photons), vertices and b-tagged jets. The HLT can therefore make a more refined accept-reject decision compared to the L1T.

Just like the L1T, the HLT has a menu of trigger seeds defining a wide range of physics signatures. The reconstruction needed to evaluate a given HLT seed, is run in steps of increasing complexity. Filter modules in-between the reconstruction steps serve to prevent waste of CPU time on subsequent steps if the event already fails certain selection criteria. This saves computational resources and contributes to more efficient trigger operation.

Events that are selected by any of the HLT trigger seeds in the menu, are accepted for offline analysis. The accepted events are sorted into HLT streams (such as *MET* or *SingleElectron*, not mutually exclusive) based on the type of HLT triggers that were fired, and permanently stored.

The LHC and the CMS Experiment

3.4 Beyond the LHC: Phase 2

By the end of Run 2 (2016-2018), the LHC delivered approximately 190 fb^{-1} of pp collision data to the CMS detector at center-of-mass energies between 7 and 13 TeV. This invaluable dataset has lead to first-time observations ranging from the Higgs boson to four-top production, as well as high precision measurements of SM processes and numerous searches for new physics realizations. At the time of writing, the LHC continues operations at twice the nominal luminosity (ie. $2 \times 10^{34} \text{ cm}^2 \text{s}^{-1}$), resulting in roughly 300 fb^{-1} of 13.6 TeV pp collision data expected by the end of Run 3 (2022-2025).

The discovery and characterization of the SM Higgs boson have played an important role in the success of the CMS physics program. However, major fundamental problems, such as the naturalness of the Higgs boson or the origin of Dark Matter, remain unanswered and demand the presence of new physics phenomena. Many models of new physics aim to solve these fundamental problems by predicting the presence of new particle states at the electroweak scale, which can be targeted via direct searches at the LHC. However, the majority of the searches already performed, including those presented in Chapters 5 and 6, show a sensitivity that is capped due to limited data statistics. Alternatively, new physics may also be exposed with precision measurements, through deviations of eg. the predicted SM Higgs coupling to charm quarks or muons, but these too are as-of-yet



Figure 3.14: Expected timeline for the LHC and HL-LHC operations [143].

unobserved due to insufficient statistical precision.

For the above measurements to sufficiently decrease statistical uncertainties, much more pp collision data is needed. However, the LHC is already running at twice its design luminosity and detectors are reaching their limits in terms of radiation damage. To continue facilitating the current analyses also after Run 3 and further increase the physics potential of the LHC, the accelerator complex will be upgraded to the High Luminosity LHC (HL-LHC) [143], building upon the experience and technology accumulated at CERN. This marks the beginning of a new era, Phase 2, with data-taking at much higher instantaneous luminosity. Consequently, also the experiments themselves will need major upgrades to continue safe and efficient data-taking. The upgrades for the HL-LHC and the CMS detector are described in Sections 3.4.1 and 3.4.2, respectively.

3.4.1 The High Luminosity LHC

The HL-LHC [143] is designed to deliver 14 TeV pp collisions at 5 to 7.5 times the nominal LHC luminosity, ie. up to $7.5 \times 10^{34} \text{ cm}^2 \text{s}^{-1}$. The expected timeline for the LHC and HL-LHC run periods is shown in Figure 3.14. LHC Run 3 will finish by the end of 2025 and, after a 3-year period of installation and commissioning, the HL-LHC will start operations in early 2029. Over the course of roughly 10 years, the HL-LHC is expected to deliver approximately 3000 fb^{-1} of pp collision data, a factor ten increase with respect to the data collected during Runs 1, 2 and 3 combined.

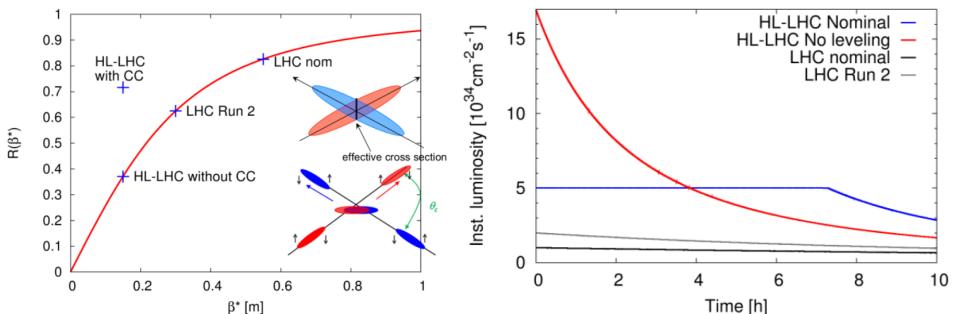


Figure 3.15: Left: Geometric luminosity reduction factor as function of β^* , indicating also the beam configurations of the LHC and HL-LHC with and without CCs. A schematic of the proton bunch rotation and overlap due to the use of CCs is included as well. Right: Instantaneous luminosity as function of time for the HL-LHC with and without luminosity leveling. Both figures from Ref. [143].

The LHC and the CMS Experiment

The increase of instantaneous luminosity at the HL-LHC is achieved in several ways. Firstly, the number of protons per circulating bunch will be increased by a factor of 2. Secondly, β^* (a measure for the beam size), will be reduced by means of more powerful triplet quadrupole magnets. Both changes to the beam configuration contribute to higher proton density, which results in more proton collisions per bunch crossing. The reduction of β^* however, implies an increased beam crossing angle θ_C , which in turn decreases the instantaneous luminosity via the geometrical reduction factor R (defined in Eqn. 3.2 of Section 3.1), hence counteracting the effect of decreased β^* . The variation of the geometrical reduction factor as function of β^* is shown in Figure 3.15 (left).

To restore the geometric reduction factor, while maintaining lower β^* and higher θ_C , a series of superconducting RF crab cavities (CCs) will be employed. Via transverse electric fields, these are able to rotate the proton bunches, such that the bunches from opposing beams perfectly overlap at the interaction point. The tilting of proton bunches, as well as its effect on the geometric reduction factor is shown in Figure 3.15 (left). As a result of the higher proton density and new bunch crossing scheme, the HL-LHC could deliver a peak luminosity up to $2 \times 10^{35} \text{ cm}^2 \text{s}^{-1}$. This corresponds to interaction rates and radiation levels that exceed the operation capabilities of current and envisioned detector technologies as well as the interaction region magnets themselves. Furthermore, as is shown in Figure 3.15 (right), the peak luminosity at the beginning of a beam fill is typically followed by a drop due to *luminosity burn-off* (proton consumption during collision). The installation of the CCs may therefore serve an additional purpose: *luminosity leveling*. Via dynamic control of the bunch tilting the instantaneous luminosity can be maintained at a constant level that is lower than the

Parameter	Symbol	LHC	HL-LHC
Center-of-mass energy [TeV]	\sqrt{s}	14	
Peak (levelled) instantaneous luminosity [$10^{34} \text{ cm}^{-2} \text{s}^{-1}$]	L_{inst}	1	17 (5)
Number of particles per bunch [10^{11}]	N_b	1.15	2.2
Number of bunches per beam	n_b	2808	2748
Events per bunch crossing (with levelling)	μ	19 (-)	212 (131)
Revolution frequency [Hz]	f_{rev}		11245
Longitudinal RMS bunch size [cm]	σ_z	7.55	9.0
Normalized transverse emittance [μmrad]	ϵ_n	3.75	2.5
β -function [m]	β^*	0.55	0.15
Crossing angle [μrad]	θ_c	285	500

Table 3.1: Design values of (HL-)LHC proton beam parameters [111, 143].

peak luminosity, but results in the same integrated luminosity as without luminosity leveling, assuming sufficient operational run-time. The flexibility to tune the luminosity and number of pile-up interactions will be an important handle for operations of the HL-LHC and the experiments, especially at the early stages of Phase 2. An overview of the relevant HL-LHC beam configuration parameters, and a comparison to the LHC parameters, is shown in Table 3.1.

3.4.2 The Phase 2 Upgrade of the CMS Detector

The record breaking instantaneous luminosity expected during Phase 2 will create an extremely challenging data taking environment; the detectors must be able to sustain the high pileup and radiation conditions. At $5 \times 10^{34} \text{ cm}^2 \text{s}^{-1}$, the average number of pileup interactions will be roughly 140, and it may reach up to 200 interactions per bunch crossing during later stages of Phase 2, ie. a factor 5 increase with respect to the 2018 LHC run period. The tremendous particle flux imposes new challenges on the CMS detector in terms of radiation hardness, energy and momentum measurement resolution, particle identification and online event selection. As such, major upgrades are being prepared [144], of which the main ones are summarised below.

Tracker upgrades

The Tracker is located closest to the interaction point and by the end of Run 3 will have already suffered significant radiation damage. It will therefore be

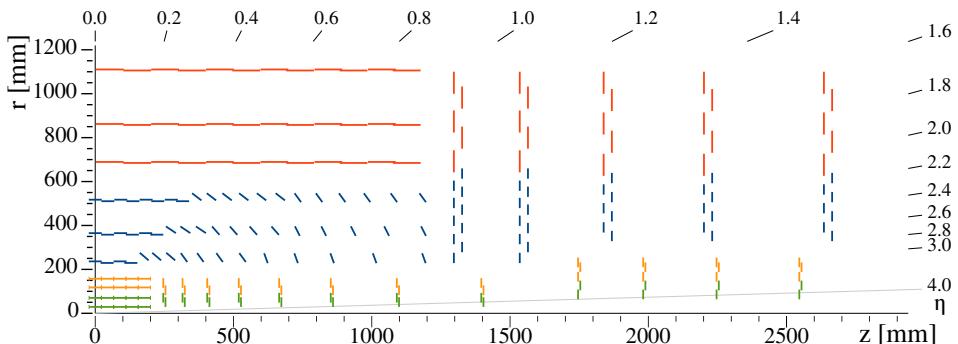


Figure 3.16: Schematic layout of one quarter of the Phase-2 CMS Tracker in the r - z plane [121]. The green (orange) segments represent pixel modules with two (four) readout chips. The red (blue) segments correspond to p_T modules with two strip sensors (a pixel and a strip sensor).

The LHC and the CMS Experiment

completely replaced by the Phase 2 CMS Tracker [121], which has been designed to operate efficiently for an integrated luminosity of up to 3000 fb^{-1} with 50% margin, and provides several improvements with respect to the current Tracker. Firstly, the silicon strips in the Outer Tracker will be shortened and the pixels in the Inner Tracker will have a decreased size of 25×100 or $50 \times 50 \mu\text{m}^2$, leading to a factor 4 increased granularity in the full Tracker volume, and sufficiently low occupancy (even at 200 pile-up). Additionally, up to 10 extra pixel disks will be installed in the forward regions, to allow efficient tracking up to $|\eta| < 4$, as shown in Figure 3.16. The reduced amount of material in the Tracker contributes to an improved p_T resolution and lower rates of γ conversion, and also benefits the performance of the Calorimeters. Last but not least, the front-end electronics of the Outer Tracker are equipped with p_T *modules*, two close-by single-sided sensors, read out by a common set of application-specific integrated circuits. They will perform local reconstruction of Tracker *stubs* at 40 MHz, which will be available for reconstruction of L1 tracks with $p_T > 2 \text{ GeV}$ by the upgraded L1T system. More details of the stub reconstruction and inclusion in the Phase 2 L1T are given in Chapter 7.

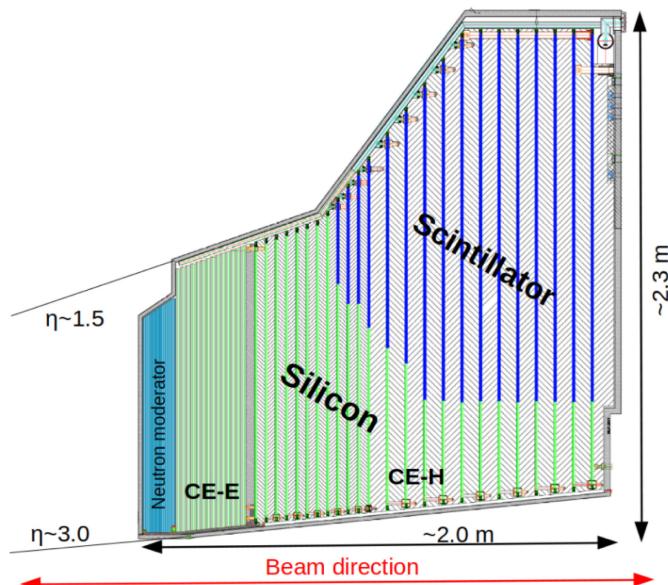


Figure 3.17: Illustration of one quarter of the Phase-2 CMS High Granularity Calorimeter (HGCAL) in the r-z plane [145].

Calorimeter upgrades

The end-caps of the ECAL (EE) and HCAL (HE, HF) will be completely replaced as well, due to the radiation damage suffered by the end of Run 3. A new High Granularity Calorimeter (HGCAL) [145] will be installed, containing both electromagnetic (CE-E) and hadronic (CE-H) sections, as is shown in Figure 3.17. The CE-E consists of tungsten and copper plates, interleaved with silicon sensors. The CE-H is equipped with brass and copper plates interleaved with silicon sensors as well, but also has a scintillator section at the back. The latter is similar to the current HE and consists of brass plates interleaved with plastic scintillating tiles. The HGCAL is designed to provide fine lateral *and* longitudinal granularity, resulting in detailed 3D descriptions of the shower shapes. Timing information will be used to aid in the rejection of pileup by associating energy deposits to interaction vertices. The barrel calorimeters will not be replaced, but receive upgrades to the detector readouts [146], primarily motivated by the L1T constraints, to generate trigger primitives with the granularity of a single EB crystal. Therefore, both the barrel and end-cap calorimeters will provide high granularity information to the Phase 2 L1T.

Muon spectrometer upgrades

Besides improved readout and electronics in the barrel region to aid in the L1T decision, the Muon spectrometer receives also major upgrades in the end-cap region [147]. In particular, new Gas Electron Multiplier (GEM) chambers will be employed in the first (GE1/1) and second (GE2/1) disk, covering $1.6 < |\eta| < 2.15$ and $1.6 < |\eta| < 2.4$, respectively, as shown in Figure 3.18. Both the GE1/1 and GE2/1 will contain two layers of GEM chambers per station. Similar to the other three muon detector technologies (DT, CSC, RPC), the GEM chambers contain a gas mixture that ionizes with incoming muons, and the resulting charge distribution on the readout electrode from the electron avalanche can be used to infer coordinates of the muon. The chambers contain three layers of $50\ \mu\text{m}$ thin insulating polyimide with copper coating on both sides of each layer. The layers have many microscopic holes edged in a hexagonal pattern, and the potential difference between the copper coatings leads to a high electric field density in the holes and a $\mathcal{O}(1000)$ gain of electric signal over the full chamber. This technology is suitable for operation at high hit rates and in high magnetic fields at the inner parts of the detector. Besides the GE1/1 and GE2/1, a third GEM station (ME0) will be installed to cover $2.0 < |\eta| < 2.8$, containing six GEM chambers per station. A

The LHC and the CMS Experiment

total of 432 GEM chambers will be installed, each providing measurements with a spatial resolution between 160 and 410 μm . Additionally, 13824 improved Resistive Plate Chambers (iRPC) will be installed in the third and fourth disk, covering $1.8 < |\eta| < 2.4$, that are able to sustain higher hit rates compared to the current RPCs, and complement the GEM and CSC detectors with an increased timing resolution of 1.5 ns.

Trigger and data acquisition upgrades

Possibly the largest challenge posed by the Phase 2 running conditions falls upon the Level-1 Trigger system. Simulation studies predict that the L1T algorithms of today in HL-LHC conditions would result in 4000 kHz of trigger rate, far beyond technical limitations [148]. The L1T will therefore be completely replaced to facilitate online event selection at 40 MHz with 200 pileup. The new, modular architecture will be based on state-of-the-art FPGA processors with increased inter-connectivity. For the first time in CMS, Tracker information will be available at 40 MHz and the L1T will receive Calorimeter information with increased gran-

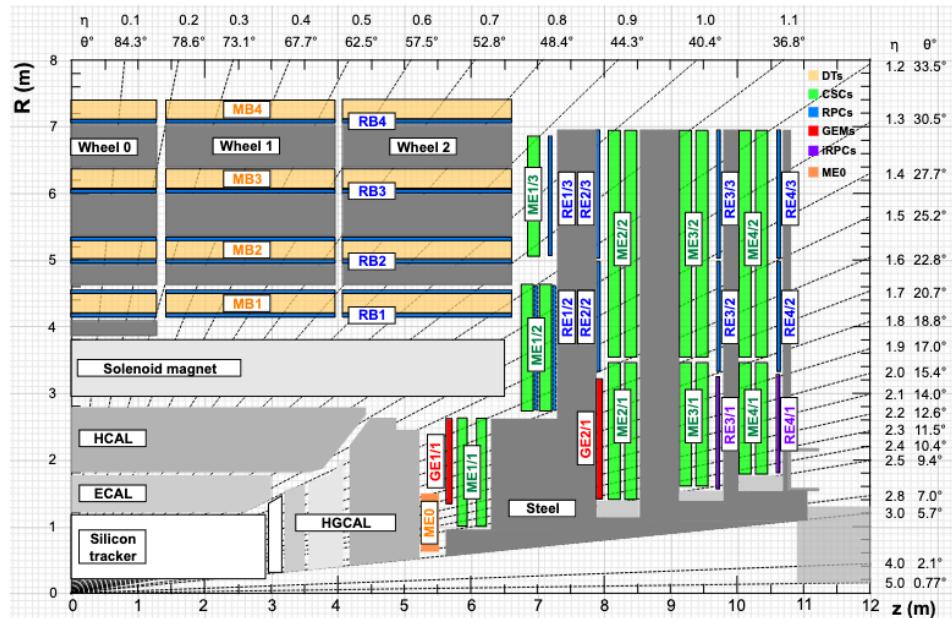


Figure 3.18: Layout of one quarter of the Phase-2 CMS Muon spectrometer in the r-z plane [147].

ularity due to the installation of the HGCAL and improved barrel Calorimeter readouts. Also the latency and maximum L1T rate are increased from $3.8\ \mu\text{s}$ and 100 kHz to $12.5\ \mu\text{s}$ and 750 kHz, respectively. The upgrades facilitate more sophisticated trigger algorithms - including machine learning techniques - that allow improved object and event-level reconstruction and aid in pileup rejection. The L1T upgrade will be described in more detail in Chapter 7. Due to the installation of new sub-detectors, in particular the HGCAL and Tracker, the total event size increases by approximately a factor of 5. The Data Acquisition (DAQ) system will therefore also be upgraded with higher bandwidths and processing power to accommodate the increased L1T rate and event size.

Chapter 4

Object and Event Reconstruction in CMS

The full reconstruction of physics objects and global quantities in a given collision event can be a truly daunting task, in particular given the hadronic, high pile-up environment of the LHC. It relies on a high detector granularity and an accurate calibration of the experiment in order to account for eg. pile-up contributions, electronics noise, detector mis-alignments and response, material interactions, et cetera.

In CMS the reconstruction is performed with the *Particle Flow* algorithm. It aims to unambiguously reconstruct and identify all particles in a given event by optimally combining information from all CMS sub-systems. The algorithm itself is introduced in Section 4.1. Next, Sections 4.2-4.4 describe the reconstruction of each type of physics object, following the order as done in Particle Flow: muons, e/γ objects and hadrons. Higher-level objects, such as jets and energy sums, are described in Section 4.5 and a dedicated pile-up mitigation technique is shown in Section 4.6.

An emphasis is put on the reconstruction of electrons, particularly at low p_T , since these are central objects in both the analysis work presented in Chapter 5 and in the Phase-2 trigger development shown in Chapter 7.

Object and Event Reconstruction in CMS

4.1 Particle Flow

At every LHC bunch crossing, a plethora of particles emerge from the proton-proton (pp) interaction region and traverse the detector volumes. They may encounter - in order - the Tracker, ECAL, HCAL and Muon sub-detectors, leaving a characteristic signature that is used for particle reconstruction and identification (see Figure 4.1). Particles carrying electric charge ionize the silicon tracker material, resulting in a series of hits along their trajectory (track), which is bent due to the magnetic field. Neutral particles generally do not interact with the Tracker, nor do they experience the Lorentz force. Electrons and photons initiate electromagnetic showers in the ECAL, while hadrons mostly shower in the HCAL. Muons interact very little with detector material and are one of the few particle types to reach the muon spectrometer, leaving hits due to ionization of the gaseous sub-detectors. Only neutrinos and hypothetical stable and weakly interacting particles, such as the lightest sparticle in many SUSY models, may escape the detector unseen.

The Particle Flow (PF) [149] algorithm aims to unambiguously reconstruct and

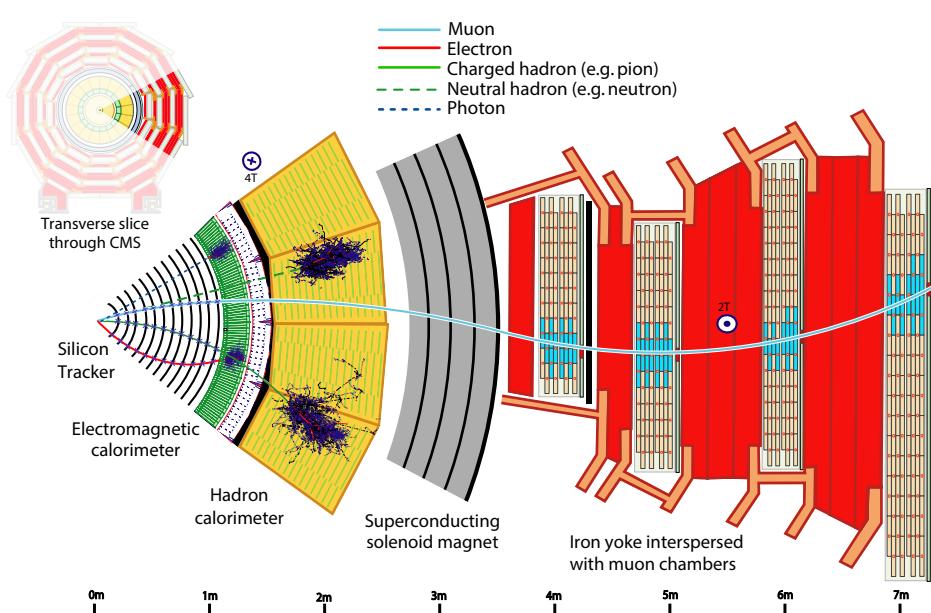


Figure 4.1: Schematic of a slice of the CMS detector in the transverse plane, indicating the characteristic signatures for particles reconstructed by Particle Flow [149].

identify all particles in a given event. It is known for its holistic event interpretation that optimally combines information from all CMS sub-systems to find the various particle signatures. The algorithm starts with the reconstruction of *PF elements*, being Tracker tracks and Calorimeter clusters, which will be described in Section 4.1.1 and 4.1.2, respectively. Compatible PF elements and hits in the Muon spectrometer may then be *linked* to form *PF blocks* of associated elements (Section 4.1.3). The PF algorithm aims to reconstruct and identify one or more physics objects from each PF block in a specific order; muons first, followed by electrons and isolated photons, and finally hadrons and non-isolated photons, as described in Sections 4.2-4.4. Once a particle is identified, the associated PF elements are removed from the PF block, and only the remaining elements may be used for the subsequent reconstruction of other particles. This resolves any ambiguities between reconstructed particles and avoids potential double counting of eg. energy deposits in the Calorimeters. Finally, the PF physics objects may also be used for the reconstruction of higher-level objects, including hadronic taus, jets and energy sums. This is described in Section 4.5.

4.1.1 KF Tracks and Vertices

The reconstruction of charged particle tracks and interaction vertices is a complex problem, involving a high level of combinatorics, in particular in high pile-up events. At the nominal LHC luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with 20 pile-up interactions¹ roughly 1000 charged particles traverse the Tracker volume [150]. The high granularity and correspondingly low occupancy of the Tracker are crucial to resolve ambiguities from closeby or partially overlapping tracks. Furthermore, a clever and robust algorithm is needed to obtain high track reconstruction efficiency and low misreconstruction rate, without exceeding the allowed computational resources.

The standard track reconstruction - widely used in CMS - is performed with a combinatorial track finder based on the Kalman Filter (KF) algorithm [151–153], henceforth referred to as *KF tracking*. It reconstructs tracks in three stages, referred to as *seeding*, *track finding* and *track fitting*, as described below. This results in a collection of Tracker tracks that is used by the PF algorithm for the identification of muons, hadrons and higher-level objects. Only for electrons a modified version of the KF tracking is employed that accounts for photon radiation (bremsstrahlung) due to interactions with the Tracker material. This dedicated electron track reconstruction is described in Section 4.3.2.

¹Note that during LHC Run 2 the average pile-up was almost double this number, as shown in Figure 3.2.

Seeding

The track reconstruction starts with the generation of seeds, small subsets of Tracker hits compatible with a charged particle trajectory. Since the trajectory of a charged particle in a quasi-uniform magnetic field is a helical path, it can be described with 5 parameters. The minimum requirement for extracting these 5 parameters is therefore 3 measurements of 3-D space coordinates anywhere in the Tracker. Seeds are generally formed with hits in the innermost Tracker layers and other hits compatible with the same track are then found by extrapolating the trajectory outwards based on the initial parameter estimates of the seed. This inside-out reconstruction is motivated by the high-granularity 3-D measurements from the pixels, that also have a lower occupancy than the strips (despite higher track density). Furthermore, it increases track reconstruction efficiency for low-momentum particles that are deflected before reaching the outer Tracker region, and is more resilient against bremsstrahlung effects [150]. Seeds are typically generated by triplets of Tracker hits, or two Tracker hits plus a constraint on the particle origin, being either a pixel vertex or the center of the beam spot. Both the pixel vertices and beam spots are formed with an extremely fast reconstruction based only on pixel hits [150, 154], that is ran before the seeding step and is also used in the HLT. It should also be noted that the Phase 1 pixel Tracker upgrade [122, 123] facilitated excellent four-hit coverage up to $|\eta| < 2.5$, allowing the generation of seeds consisting of hit quadruplets as well.

Track finding

The second stage of the track reconstruction is based on the Kalman filter method that starts from a given seed and aims to progressively build the full track by extrapolating the trajectory and collecting hits from successive detector layers [150]. All material of the Tracker is assumed to be within the detector layers, such that the propagation *between* detector layers can be performed with a fast analytical propagator following a helical path. The propagation *inside* a detector layer is done with a more accurate material propagator that also takes into account effects from material interactions. More specifically, multiple Coulomb scattering is treated as a stochastic process with Gaussian scattering angle distribution and its predicted width is used to increase the track parameter uncertainties. Furthermore, the track momentum is reduced by the average expected Bethe-Bloch like

energy loss [155], which is dominated by ionization². The corresponding uncertainty is best modeled as a Landau probability density function (PDF), but can be approximated with a Gaussian PDF instead (ignoring the extended tail towards higher energies) to match the KF formalism. After propagating the trajectory to the next Tracker layer, compatible hits within that layer are found via a χ^2 test, incorporating uncertainties both from the hits and the trajectory itself. New track candidates are formed for each compatible hit by adding the hit to the initial candidate and updating the track parameters and covariance matrix. To limit computational requirements, a maximum of 5 new track candidates are retained per layer per initial candidate. In case more candidates are formed, the 5 best ones are chosen based on the χ^2 and total number of hits. All new track candidates are then propagated to the next layer, and the procedure is repeated until a stopping condition is reached, which typically consists of a track reaching the Calorimeter surface or containing too many missing hits.

Track fitting

The third and final stage aims to find the best estimates of the track parameters at any point along the trajectory, and reduce bias from eg. beam spot or pixel vertex constraints imposed in the seeding step [150]. For every candidate track a Kalman filter is initialized based on the inner-most hits that iterates outward, progressively updating the track parameters and uncertainties with every hit. After this filter, a smoothing step follows, which consists of a second Kalman filter that is run backwards (ie. outside-in). Therefore, at any given point along the trajectory, the first filter estimates the track parameters based on hits found *before* the reference point, while the estimates from the second filter are based on hits found *after* the reference point. The final trajectory parameters are obtained from a weighted average of the estimates from both filters. For optimal precision, the propagation of the state vectors in both filters is performed with a Runge-Kutta propagator that takes not only material interactions into account, but also inhomogeneities in the magnetic field. After the final fit, selections are applied on the number of hits, fit quality and vertex compatibility, to significantly reduce the fraction of fake tracks.

²This works well for all charged particles except electrons, which instead follow a Bethe-Heitler like energy loss, dominated by bremsstrahlung, with non-Gaussian probability density function [156, 157].

Object and Event Reconstruction in CMS

Iterative tracking

Instead of performing a single iteration of the above KF tracking, a higher efficiency for the same misreconstruction rate can be achieved by performing multiple successive iterations, each with different seeding configurations and quality criteria, targeting varying kinematic topologies. After each iteration the hits corresponding to the reconstructed tracks are masked, such that less hits are available for the next iteration. This allows successive iterations to become increasingly complex, increasing reconstruction efficiency, while staying within computational resource limitations. During 2016 data-taking, 10 KF tracking iterations were used [149]. The first iterations mainly target prompt, high quality tracks seeded by three pixel hits, which allowed reconstruction down to $p_T=200$ MeV. The next iterations aim to reconstruct tracks with one or two missing pixel hits (addressing detector inefficiencies), displaced tracks (corresponding to eg. decays of long-lived particles) and tracks in high- p_T jets. The last two iterations were newly added prior to Run 2 in order to address an observed pile-up dependent efficiency loss for muons [159], and rely also on information from the reconstruction in the Muon spectrometer (Section 4.2). The first of these two starts from a *tracker muon* and rebuilds the inner track with looser quality criteria in order to increase the track hit efficiency. The other muon-specific iteration uses *standalone muon* tracks that pass $p_T > 10$ GeV and a minimum set of quality requirements, to seed an outside-in inner track re-

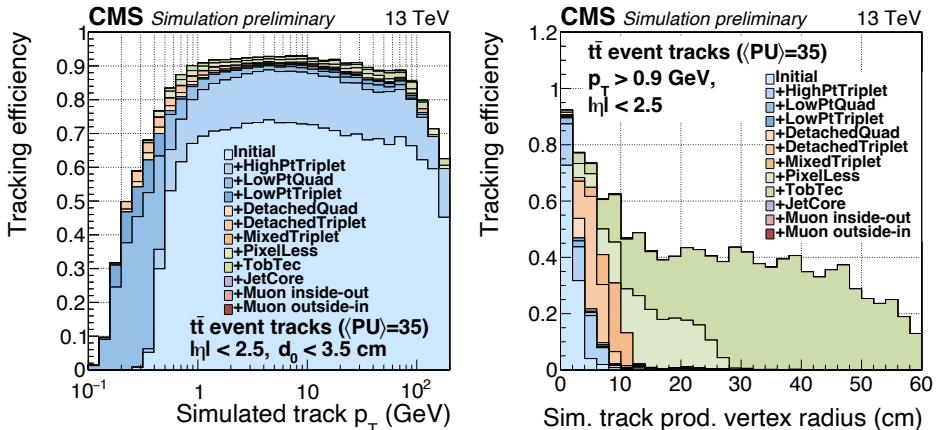


Figure 4.2: Cumulative reconstruction efficiencies for the tracking iterations employed after the Phase 1 pixel Tracker upgrade as function of track p_T (left) and transverse vertex displacement (right) [158]. The efficiencies are based on simulated $t\bar{t}$ events with 35 pile-up collisions.

construction step. Lastly, after the pixel Tracker upgrade [122, 123], the tracking in 2017 and 2018 employed 12 iterations, modifying the pixel-only iterations; those based on only hit-pairs were removed to reduce the fake rate in high pile-up events, and those based on triplets were replaced by quadruplet iterations, exploiting the excellent four-hit coverage introduced by the new pixel layers [160]. The resulting tracking efficiency in simulated $t\bar{t}$ events is shown in Figure 4.2.

Vertex reconstruction

Primary vertices are reconstructed from the final set of tracks by sequentially selecting, clustering and fitting the tracks to a common position in the primary pp interaction region. The selection is based on the compatibility of tracks with coming from prompt decays, the fit quality and number of hits in the pixel and strip detectors [150]. The clustering of tracks into a set of tracks corresponding to the same pp interaction vertex is performed with a deterministic annealing algorithm [161]. The best estimates of the vertex are then obtained by fitting the tracks from a given set using an adaptive vertex fitter [162]. During Run 1, the primary vertex from the hard-scatter was identified as the one having the highest quadratic sum of p_T of the associated tracks, $\sum p_T^2$. However, with increasing number of simultaneous pp collisions this tends to misidentify the hard-scatter vertex. As of Run 2, it is therefore identified instead by clustering the tracks associated to each vertex into track jets using the anti- k_T algorithm [140, 163], and picking the vertex with highest quadratic sum of track jet p_T [164]. The remaining vertices are classified as pileup vertices.

Secondary vertices may occur through the decay of particles with a considerable lifetime, such as hadrons (B , Λ , K_S^0 , etc.), taus or hypothetical new physics states. They can also occur through material interactions, with photon conversions and nuclear interactions of hadrons being the most relevant, as these may result in a pair of charged particle tracks that is displaced with respect to the corresponding primary vertex. On the other hand, decays of a charged particle into one charged and one or more neutral components, such as bremsstrahlung, will simply look like a kink in the track, and not as a secondary vertex per se. The reconstruction of the secondary displaced objects profits from the later iterations of track seeding, as these specifically target tracks with displacements up to 60 cm [149]. The identification of the displaced tracks and secondary vertices for each type of process is performed with dedicated algorithms (see Sections 4.3.3 and 4.5.1) [165–168].

4.1.2 Calorimeter Clusters

A well-performing procedure for the clustering of energy deposits in the Calorimeter is important for the measurement of the energy and direction of photons and neutral hadrons, as well as for charged hadrons in case the corresponding track parameters lack accuracy due to eg. low quality track fits and high p_T particles. It can provide discrimination power for various particle types, and contributes to the reconstruction of electrons by facilitating the identification of associated bremsstrahlung photons. The latter is done by grouping Calorimeter clusters located along tangents to the electron track into *superclusters*, as will be described in Section 4.3.

The baseline clustering of calorimeter deposits is performed separately for the various ECAL and HCAL components (EB, EE, ES, HB and HE), using varying seeding and clustering configurations, but following the same general strategy [149]. No clustering is performed for HF, as each cell already provides electromagnetic and hadronic clusters.

The clustering starts by finding *cluster seeds*, defined as cells with local energy maxima passing a threshold that depends on the sub-detector. The energy maxima are identified by comparing a given cell energy to that of its 4 or 8 nearest neighbours for the HCAL and ECAL, respectively. *Topological cluster* are then built iteratively by adding cells that contain energies of at least twice the noise level and that share a side or corner with a cell already part of the cluster.

Next, an assumption is made, postulating that the total energy distribution in the topological cluster is the result of a mixture of N Gaussian energy deposits, with N being the number of seeds within the topological cluster. To resolve the individual deposits, the parameters of the Gaussian functions are initialized with the energy and position of the corresponding seeds. An iterative procedure then follows, which alternates between estimating the fraction of energy in a given cell from each of the Gaussians with fixed parameters, and improving the parameter estimates with a maximum likelihood fit for fixed energy fractions. The final parameter estimates after convergence define the individual energy deposits. A schematic example of the clustering procedure is shown in Figure 4.3.

The energy thresholds imposed in the seed and topological cluster generation of the ECAL reduce the total energy of measured deposits compared to what is expected from incoming particles. A generic calibration is applied to all ECAL clusters, depending on the measured energy and pseudorapidity, accounting also for the possibility of a single particle in the forward regions to deposit energy in

both the ES and EE. The corrections are derived from simulated single photon samples and after calibration the energy matches the true photon energy within 1% over the full energy spectrum and ECAL fiducial region [149].

The energy calibration for HCAL clusters accounts for the cell energy thresholds, but also for the fraction of energy that hadrons lose through deposits in the ECAL. Simulated K_L^0 hadron samples are used to derive the corrections as function of energy and pseudorapidity, based on the calibrated ECAL energy, measured (raw) HCAL energy and true energy of simulated hadrons. After calibration, the energy response is within 1% at high energies, increasing up to 4% at $E < 20$ GeV [149].

4.1.3 PF Element Linking Scheme

As mentioned in Section 4.1, the PF algorithm employs a linking scheme to mark associated PF elements and form PF blocks, reflecting the detector signature of one or more physics objects. The capability of PF to correctly link the PF elements and identify all particles depends on the detector granularity, the particle density and the level of material interactions causing kinks in particle trajectories or creating secondary particles. The general linking strategy between any pair of elements

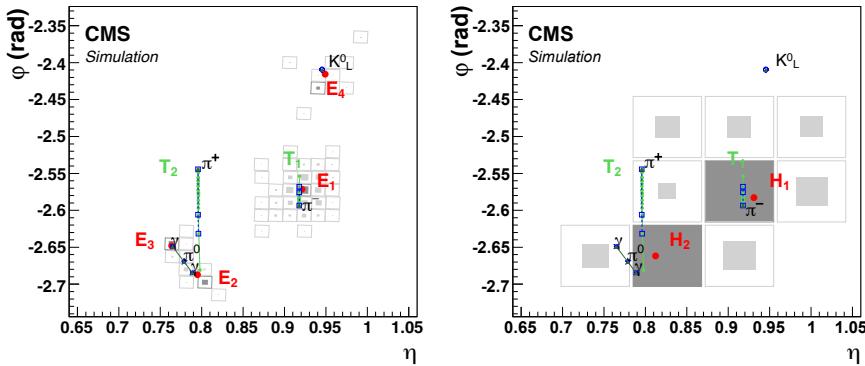


Figure 4.3: Example event display illustrating the clustering in ECAL (left) and HCAL (right) and PF linking scheme [149]. Grey squares indicate Calorimeter cells with the inner area reflecting the cell energy and local maxima marked in dark grey. The K_L^0 , π^- and photons (from $\pi^0 \rightarrow \gamma\gamma$) result in three topological ECAL clusters. The π^+ leaves no energy deposit in ECAL. The tracks ($T_{1,2}$) from the charged pions point to the HCAL deposits. The estimated positions of the individual ECAL and HCAL clusters are shown with the red dots (E_{1-4} and $H_{1,2}$, respectively). Four PF blocks are formed: 1 for the charged pions, linking T_1 , E_1 , H_1 , T_2 and H_2 , and 1 block for each of the remaining ECAL clusters of the K_L^0 and photons.

in a given event is described in the next paragraph and an example is shown in Figure 4.3.

A KF track is linked to a Calorimeter cluster if the track, extrapolated through the Calorimeter volumes, passes through any of the cells associated to the cluster. A margin of one cell around the cluster area is allowed to account for gaps between Calorimeter cells, cracks, and the uncertainty in track extrapolation and cluster shape. ECAL clusters are linked to HCAL clusters if their position is within the cluster envelope in the HCAL. The same is done for linking preshower clusters to those in the ECAL end-caps. A 2-D link distance is defined based on the cluster or extrapolated track positions. If multiple HCAL clusters are linked to the same track or ECAL cluster, or if multiple ECAL clusters are linked to the same preshower cluster, or if multiple tracks are linked to the same ECAL cluster, then only the links with the smallest distance are retained. Finally, KF tracks may be linked together if they share a common secondary vertex that corresponds to a nuclear interaction with detector material. In this case, at least three tracks must be associated to the secondary vertex, at most one of those is an incoming track, and the invariant mass of the outgoing tracks must be larger than 0.2 GeV.

Specific linking procedures are implemented for the reconstruction and identification of photons and electrons, as they are intimately connected through photon conversions and bremsstrahlung. This is described in detail in Section 4.3. Also for muons, a dedicated procedure is employed to link KF tracks to hits in the Muon spectrometer. This is described in Section 4.2 and is strictly not part of the PF algorithm itself.

After linking the elements, PF identifies particles in a specific order, starting with muons, then electrons and isolated photons, and finally hadrons and non-isolated photons. Once a particle is identified, the corresponding PF elements are *locked*, such that they cannot be used for successive particle identifications.

4.2 Muons

Muon reconstruction and identification exploits the high stopping power of the Calorimeters. By design, this contributes to a high purity of muon hits in the Muon spectrometer, with a minor background of remnants from hadrons *punching through* the Calorimeters and reaching the innermost muon stations.

The muon reconstruction [130, 169] starts with the local reconstruction of track segments in the DT and CSC sub-detectors and hits in the RPC sub-detector, based on information from only a single muon chamber. The DT and CSC track

segments consist of a small set of compatible hits and give a rough estimation of the corresponding track position, momentum and direction. The RPC hits are simply points in the detector plane. This local reconstruction is followed by the reconstruction of *muon tracks*, relying either solely on the Muon spectrometer information (segments, hits), or additionally on the KF tracks found from hits in the Tracker (Section 4.1.1). Since the muon track reconstruction itself also relies on the KF algorithm, the KF tracks found from Tracker hits only are here referred to as *inner tracks*. Three types of muon tracks are reconstructed.

Standalone muons rely solely on Muon spectrometer information. The tracks are constructed with a KF procedure, very similar as described in Section 4.1.1. The track is seeded by segments from the DT and CSC sub-detectors and built inside-out by progressively adding more measurements (DT segments or hits in the CSC or RPC) compatible with the track, updating the track parameters and uncertainties at each new measurement. After also the backward (smoothing) fit is performed, the track is furthermore fit to the center of the beam-spot under the assumption that is originates from prompt decays in the pp interaction region.

Tracker muons are reconstructed starting from all inner tracks with (transverse) momentum above (0.5) 2.5 GeV. The selected tracks are propagated to the Muon spectrometer (inside-out), while taking into account material interactions and the magnetic field. For the track to be compatible with the muon hypothesis, the energy deposits in each Calorimeter volume within a distance ΔR around the track are summed and may not exceed a certain threshold. The values of ΔR and energy thresholds are different for each Calorimeter. Furthermore, the track must be well-matched to muon segments, by passing cuts on the number of matched segments and their position with respect to the extrapolated track [169, 170]. The fact that only a single muon segment is required for Tracker muons, makes them more efficient to low momentum muons (below $p \simeq 5\text{GeV}$), since those typically only penetrate one muon station before their track is lost due to multiple scattering inside the magnet return yolk.

Global muons are reconstructed from a standalone muon track that is propagated to the Tracker (outside-in), while taking into account material interactions and the magnetic field. If a compatible inner track is found, a combined track fit is performed using the hits from both the inner track and standalone muon track, forming the global muon track. This combined fit improves the momentum resolution for muons with p_T above $\simeq 200\text{ GeV}$. Roughly 99% of muons are reconstructed as global or tracker muon, often as both. If a global muon and a tracker muon share the same inner track, they are merged into a single muon object.

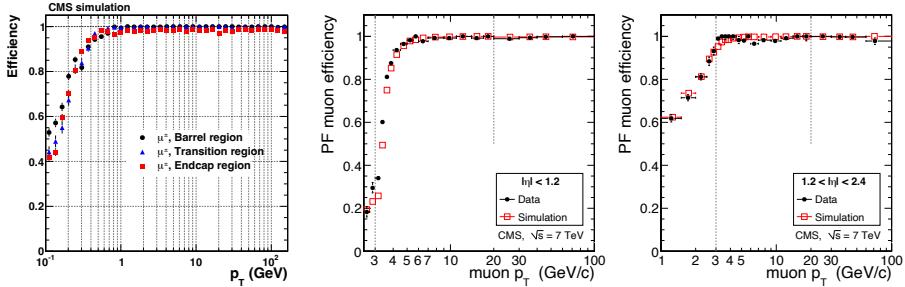


Figure 4.4: Left: The track reconstruction efficiency for muons in the barrel, transition and end-cap regions, defined by $|\eta|$ intervals of 0-0.9, 0.9-1.4 and 1.4-2.5 [150]. Middle and right: The combined muon reconstruction and PF identification efficiency, given a reconstructed inner track, shown for $|\eta|$ intervals of 0-1.4 and 1.4-2.4 [171]. Note that these figures show the Run 1 performance, and therefore do not include more recent improvements, such as the muon-specific tracking iterations added for Run 2. More recent results are shown in Refs. [130, 159], although those studies do not include the p_T dependency down to $O(1)$ GeV. Instead, the muon reconstruction and identification efficiency at low p_T is studied in more detail in the search presented in Chapter 5 (Figure 5.19).

To be identified as a *PF muon*, the reconstructed global or tracker muons must pass additional selection criteria [149, 172]. Firstly, isolated muons are selected if the sum of track p_T and calorimeter energy deposits within ΔR of the candidate do not exceed 10% of the muon p_T . Selections tuned for identification of muons inside jets are applied on the remaining candidates, and depend on the number of matching track segments, the match quality and the calorimetric footprint. The efficiency of muon track reconstruction and PF muon identification is shown in Figure 4.4. After PF muon identification, the associated elements are locked, such that they cannot be used for the identification of other particles.

4.3 Electrons and Isolated Photons

After PF has identified the muons and locked the associated PF elements, the algorithm proceeds with the reconstruction and identification of electrons and isolated photons. These objects are intimately connected through bremsstrahlung ($e \rightarrow e\gamma$) and conversion ($\gamma \rightarrow ee$) processes. Electrons that are produced in the pp interaction region with $|\eta| \simeq 0$ (1.4) on average emit 33 (86)% of their initial energy in the form of photon radiation before reaching the ECAL [173]. On the other hand, up to 60% of photons may convert before reaching the last three Tracker layers [167]. Therefore, for a given electron, bremsstrahlung and subsequent photon conversions may result in additional ECAL clusters and charged particle tracks,

which are all associated to the original electron. To accurately measure the initial electron energy, it is crucial to identify and group the associated objects from photon radiation and conversion. Although the standard Calorimeter clusters and KF tracks *are* used as input, dedicated clustering and tracking algorithms are employed for the reconstruction and identification of e/γ objects [174].

Firstly, *mustache superclusters* are built using only Calorimeter information, by grouping ECAL clusters based on their (η, ϕ) location and energy, as described in Section 4.3.1. Next, a modified version of the KF tracking algorithm, partially seeded by these superclusters, aims to reconstruct the electron tracks, as described in Section 4.3.2. Finally, the tracks and clusters are combined to *refine* the supercluster and identify photon conversions. This is described in Section 4.3.3

The refined superclusters, ECAL clusters, standard KF tracks (with and without conversion flag) and electron tracks are all used to reconstruct the Global Event Description (GED) and PF electrons, as described in Section 4.3.4. These are the standard electron objects that are commonly used in most CMS analyses. However, an alternative reconstruction was developed [175, 176], dedicated to low p_T electrons and originally motivated by analyses measuring $R_{K^{(*)}}$ in the context of flavour anomalies. This alternative reconstruction will be described in Section 4.3.5 and is also employed by the search for supersymmetry with compressed mass-spectra that is presented in Chapter 5.

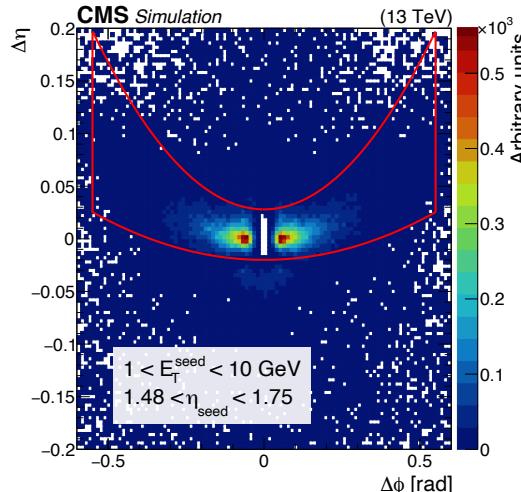


Figure 4.5: Distribution of η and ϕ of clusters with respect to a seed cluster corresponding to simulated electrons with $1 < E_T^{\text{seed}} < 10$ GeV. The red line indicates the mustache-shaped window used to build the supercluster [174].

4.3.1 Mustache Superclusters

The mustache superclustering algorithm [174] aims to group ECAL and preshower clusters that are compatible with the same initial electron object. A cluster with transverse energy E_T above a given threshold is used to seed the algorithm. Due to the magnetic field configuration, associated clusters are found mostly along the ϕ direction, while there is only a small η -dependency for low p_T objects. Therefore, a window in the (η, ϕ) plane with a mustache-like shape (hence the name) is built around the seed cluster, as shown in Figure 4.5. The size of the mustache window depends on the energy of the seed cluster, reflecting the fact that the trajectories of charged particles with higher momenta are bent less by the magnetic field. Clusters located inside the mustache window may be added to the seed cluster to form the so-called *mustache supercluster*.

4.3.2 GSF Tracking

As mentioned in Section 4.1.1, the standard Tracker tracks are reconstructed with a combinatorial track finder based on the KF algorithm. This accounts for Bethe-Bloch like energy loss of charged particles inside detector material, which is dominated by ionization [155]. Furthermore, the associated Landau uncertainty can be modeled to a good approximation by a Gaussian PDF, conform the KF formalism. This works well for all charged particles, except electrons, since the energy loss of electrons is mostly due to bremsstrahlung. The energy loss of electrons can instead be modeled with the Bethe-Heitler formula and a highly non-Gaussian PDF [156, 157]. The reconstruction of electron tracks therefore relies on a modified version of the standard KF tracking that instead is based on a Gaussian-sum filter (GSF) algorithm [150, 157, 173, 177]. It follows a similar procedure as the KF tracking described in Section 4.1.1, but with several differences in each of the seeding, track finding and track fitting steps, as summarised below.

Electron seeding

Since the GSF tracking is very computationally intensive, the seeding step aims to find small subsets of Tracker hits in the innermost layers, that are compatible with originating from an electron. Two seeding procedures are used. *ECAL-driven seeding* starts from the set of reconstructed superclusters that satisfy $E_{SC,T} > 4$ GeV and $H/E_{SC} < 0.15$, where $E_{SC,T}$ and H represent the supercluster energy and the sum of energy deposits in the HCAL towers within $\Delta R=0.15$ around the super-

cluster position, respectively [174]. For each of the selected superclusters, a state-vector is initiated based on the supercluster parameters and propagated inwards along a helical trajectory, ignoring material interactions, and for both charge hypotheses. The extrapolated trajectory is matched to hits in the innermost Tracker layers using geometrical windows, depending on the charge hypothesis and $E_{SC,T}$, in order to generate the ECAL-driven seeds. *Tracker-driven seeding* instead start from the collection of standard KF tracks with $p_T > 2$ GeV [149]. If the track corresponds to an electron that radiated little to no energy, then the KF track typically has a high number of hits, well-behaved χ^2 and forms a good track-cluster match. Therefore, if a KF track is matched to an ECAL cluster via standard PF linking (Section 4.1.3) and has $E_{\text{cluster}}/p_{\text{track}}$ close to unity, then its seed is selected for GSF tracking. On the other hand, electrons that radiate a large fraction of their energy will typically be reconstructed by KF tracking as having few hits and large χ^2 . Those tracks are fit again with a Gaussian-sum filter. A GSF is similar to a KF in the sense that it iteratively propagates the state-vector, while accounting for energy loss and the magnetic field, and adds more measurements (hits), each time updating the state-vector parameters and uncertainties. The key difference lies in the modeling of the energy loss and the associated uncertainty. The GSF reduces the track momentum by the average expected Bethe-Heitler like energy loss, while the standard KF uses the Bethe-Bloch formula. Furthermore, the uncertainty on the energy loss in the GSF is modeled with a Gaussian mixture (ie. a weighted sum of Gaussian PDFs), while the KF only uses a single Gaussian PDF. A more accurate modeling of the uncertainty can be achieved by using more weighted Gaussians in the GSF, but this also implies a large increase in computational requirements. Therefore, at this seeding stage, the refitting of the KF track with a GSF is done using only 5 Gaussians - referred to as *components* - while the full GSF track fitting performed at later stages relies on 12 components. The χ^2 of both the GSF and KF fits, the number of hits, the total energy loss and the distance between the track (extrapolated to the Calorimeter surface) and the closest ECAL cluster are used by a boosted decision tree (BDT) [178] to select the track seeds for the full GSF tracking. The ECAL-driven seeding performs better at high energies, reaching a 95% seeding efficiency for electrons with $E_T > 10$ GeV from Z boson decays. The Tracker-driven seeding recovers efficiency for softer or non-isolated electrons, having a seeding efficiency above 50% at $p_T > 3$ GeV [174]. Both seed collections are merged and used to initiate the track finding as described below.

GSF track finding

The electron track finding is nearly identical to the standard track finding described in Section 4.1.1. It is based on a standard KF to progressively build the electron track, but instead relies on the Bethe-Heitler formula to model energy loss. The associated uncertainty is modeled with a single Gaussian (ie. *not* with a Gaussian mixture), but the compatibility between the projected trajectory and Tracker hits is determined with an increased χ^2 threshold. A penalty is applied for missing hits by increasing the track χ^2 in order to prevent picking up hits from converted photons, and at most one missing hit is allowed along a given trajectory.

GSF track fitting

Just like in the standard KF tracking, a final fit is performed after the track is built, to determine the final track parameters. The electron track fit is based on a GSF with 12 components to optimally account for the energy loss and associated uncertainty. The forward GSF fit is followed by a backward (smoothing) GSF fit, and the final track parameters at any point along the trajectory are given by the weighted average of both fits.

4

4.3.3 Refined Superclusters and Conversion Finding

Bremsstrahlung photons are emitted due to material interactions, and therefore occur at the Tracker layers traversed by the electron. Furthermore, the angular distribution of the emitted photons is driven by the Lorentz factor $\gamma_e = E_e/m_e$, where E_e and m_e are the energy and rest-mass of the electron. The average emission angle is $\theta \simeq 1/\gamma_e$, which at typical LHC energy scales is close to zero [179]. Secondary objects from bremsstrahlung and subsequent conversions may therefore be found along tangents with a given GSF track, generated at each of the traversed Tracker layers.

Unconverted bremsstrahlung photons are identified if a generated tangent along the GSF track directly points to an ECAL cluster. In that case, a link is created between the GSF track and the ECAL cluster, allowing recovery of the bremsstrahlung photons if the cluster was not already part of the mustache supercluster.

Converted bremsstrahlung (and primary) photons are more challenging to reconstruct, due to lower resolution of displaced tracks, large combinatorics and the possibility of additional bremsstrahlung after conversion. The general reconstruction strategy is to search for oppositely charged track pairs compatible with a

common vertex [165–167]. To maximize efficiency, a highly inclusive set of tracks is used, consisting of both standard KF tracks and GSF tracks. Furthermore, an ECAL-driven method is employed to reconstruct tracks from late-occurring conversions in the outermost Tracker layers [180, 181]. In this method, trajectory state-vectors are initiated based on the constituents of superclusters and nearby ECAL clusters and propagated inwards (similar to the ECAL-driven seeding for GSF tracking). If compatible hits in the three outermost Tracker layers are found, then the track building continues inwards with a KF accounting for Bethe-Heitler like energy loss. The resulting tracks are then considered to be one arm of the electron-positron pair, and the innermost hits are used to seed outward track finding with the aim of reconstructing the other arm. This method recovers efficiency for high p_T photons and conversions far from the pp interaction regions, but is not suitable for reconstruction soft photon conversions, as those electrons are less likely to reach the ECAL. The inclusive track set is constructed by merging the collection of standard KF tracks, GSF tracks and the ECAL-seeded tracks. To avoid double counting tracks reconstructed from the same hits, duplicates are typically identified based on hits shared between respective tracks, and resolved by keeping only the track with highest χ^2 or largest number of hits. The inclusive track collection is used to search for pairs with conversion-specific signatures. For example, the tracks must have opposite charge, be compatible with a common vertex (constructed with a 3-D vertex fit of the selected tracks), where their trajectories are approximately parallel, and satisfy additional kinematic cuts. Besides the identification of conversion track pairs, an alternative method employs a BDT to identify tracks from photon conversions in case only one leg is reconstructed [174]. If a conversion is identified, then the tracks may be linked to the compatible ECAL clusters, and the direction of the reconstructed photon momentum may be used to link the conversion tracks to the original GSF track [149].

4.3.4 GED and PF e/γ Objects

The reconstruction and identification of electrons and isolated photons is based on mustache superclusters, ECAL clusters, the standard KF tracks, GSF tracks and conversion-flagged KF tracks (such as those reconstructed by the ECAL-driven conversion finder method). The objects are linked into PF blocks, from which one or more e/γ objects are built. The resulting Global Event Description (GED) photons and electrons are the objects commonly used by CMS analyses. Although these are sometimes referred to as *PF-like*, they are not the same as PF photons

and electrons, as explained below.

GED photons are constructed from superclusters, refined with compatible ECAL clusters and associated tracks, and are required to pass a loose selection on the supercluster energy ($E_{SC,T} > 10$ GeV) and H/E , where H represents the sum of energy deposits in the HCAL towers within $\Delta R = 0.15$ around the supercluster position [149]. The supercluster position is defined as the energy-weighted average of the positions of the supercluster constituents.

GED electrons are constructed by first matching a GSF track to a supercluster. For both ECAL-driven and Tracker-driven GSF tracks, a BDT is employed that combines kinematic and quality related variables of the track, shower shape variables of the supercluster and track-cluster matching variables. ECAL-driven GSF tracks may also be matched to a cluster via simple ($\Delta\eta, \Delta\phi$) cuts, comparing the position of the track extrapolated to the Calorimeter with the supercluster position [174]. All associated ECAL clusters and tracks are used to refine the electron candidate, and a minimal selection is applied to qualify as GED electron.

A correction to the total energy of the collected ECAL clusters serves to account for any object that was missed in the association procedure. The energy correction is assigned to the photons and can be as high as 25% depending on the material budget [149]. The final electron energy is based on the corrected ECAL energy as well as the GSF track momentum, while the electron direction is dictated fully by the GSF track.

The identification of GED photons and electrons is greedy by design, resulting in high object efficiency, but also a high fake-rate. Therefore, additional identification criteria are to be applied at the analysis level. It should also be noted that a given object may be both a GED photon and a GED electron. This differs from the PF algorithm, which aims to unambiguously reconstruct all physics objects. To qualify as a PF e/γ object, the GED e/γ must therefore pass additional selection criteria [149].

PF electrons are a strict subset of the GED electron collection. The additional selection criteria include a multivariate selection based on BDTs that are trained separately in the barrel and endcap region, as well as for isolated and non-isolated electrons. Input features are based on eg. the radiated energy fraction, fit qualities, number of hits and geometrical configurations. The PF electron efficiency for low p_T electrons is shown in Figure 4.6 (right).

PF photons are selected from the GED photons based on isolation criteria and the compatibility of the energy deposits in ECAL and HCAL with the photon hypothesis. PF photons are not a strict subset of GED photons, since the PF photon

collection may also include non-isolated photons, as described in Section 4.4.

After identifying PF e/γ objects, the associated PF clusters and tracks are locked to avoid using them for identification of other particles. It should be noted that e/γ identification relies on GSF tracks, while all other physics objects rely on the standard KF tracks. The PF elements to be locked are therefore the KF tracks, specifically those that are either already associated to the electron object (via eg. the tracker-driven GSF seeding step), or by matching to a GSF track based on shared hits or geometrical criteria. Lastly, the components of GED e/γ objects that fail the PF selection may be used by PF for the reconstruction of other particles, such as hadrons or non-isolated photons.

4.3.5 LowPt Electrons

The GED (and PF) electrons described in Section 4.3.4 are the standard electron objects used by most CMS analyses. However, a new and independent *LowPt electron* reconstruction algorithm was designed with the aim of increasing efficiency to softer objects [175,176]. It was originally developed for the 2018 B-parking data-stream, which enabled the accumulation of 10 billion unbiased B hadron decays,

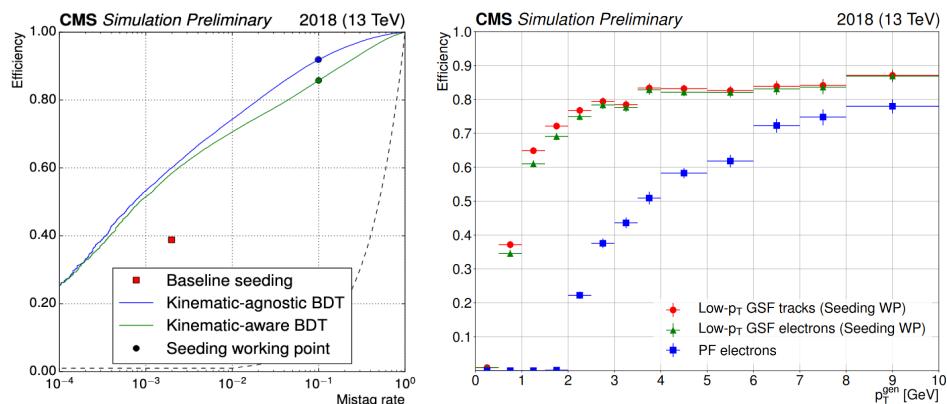


Figure 4.6: Left: ROC curves for the two BDT based models that are implemented in a new Tracker driven GSF seeding logic designed for LowPt electrons. The ROCs are evaluated on $B^+ \rightarrow K^+ e^+ e^-$ events. The indicated seeding working points correspond the loose working points applied in the LowPt electron reconstruction of the B-parking data-set. Tight working points are also defined (not indicated) for the LowPt electron reconstruction in the standard data-sets, corresponding to a 0.2% mistag rate. These result in a 50% gain of efficiency at the same mistag rate with respect to the baseline GSF seeding used for GED electrons (red square). Right: Reconstruction efficiencies of GSF tracks with the LowPt seeding logic (red) and LowPt electrons (green) in the B-parking data-set, compared to the standard PF electrons (blue). The LowPt electron efficiencies for the standard data-set are shown in Chapter 5 (Figure 5.15). [175]

in order to provide measurements of R_K and R_{K^*} competitive with those of the BaBar, Belle and LHCb experiments³. This LowPt electron reconstruction has furthermore been adopted in the re-reconstruction of the standard full Run 2 CMS data-sets.

Similar to the standard GED electrons, the reconstruction of LowPt electrons relies on a GSF track that is associated to a cluster in the Calorimeter. The key difference lies in the GSF seeding. Instead of the usual ECAL- or Tracker driven seeding, the LowPt electron reconstruction employs a new Tracker driven seeding logic that relies on two independent models based on BDTs; One model aims to identify electron seeds using as input features several track parameters, track-to-cluster matching features and shower shape variables, and is referred to as the “unbiased” or “kinematic-agnostic” BDT. The other model relies on the same input features, as well as the track p_T , η and d_{xy} significance, hence called “ p_T -biased” or “kinematic-aware” BDT. For the LowPt electron reconstruction in the B-parking data-set of CMS, loose working points are defined by evaluating the ROC curves (obtained from $B^+ \rightarrow K^+ e^+ e^-$ events) and correspond to a mistag rate of 10%, as is shown in Figure 4.6 (left). KF tracks that pass the corresponding thresholds of either BDT are used to seed the GSF track reconstruction for objects with p_T down to 0.5 GeV. The efficiencies of the GSF tracking with the new seeding logic and the LowPt electron reconstruction are shown in Figure 4.6 (right). A greatly improved efficiency is obtained for soft electrons with respect to the standard PF electron reconstruction. The LowPt electron reconstruction is also performed for the standard data-sets, but with slightly tighter requirements⁴. Tighter seeding working points are used, corresponding to a mistag rate of 0.2% and the p_T cut is increased to 1.0 GeV. Furthermore, LowPt electrons stored in the MiniAOD data-tier of the standard data-sets are required to pass another MVA ID (2020Nov28), that aims to further reduce the amount of fake LowPt electrons. LowPt electron reconstruction is independent of PF, hence no PF elements are locked once a LowPt electron is identified. As such, the seeding KF track and Calorimeter cluster of the LowPt electron can still be used by PF for the reconstruction of other objects. To avoid double-counting objects that are reconstructed by both algorithms, a custom overlap removal step has to be used, such as presented in Chapter 5 (Section 5.5.1).

³This was mainly motivated by tests of lepton flavor universality by the aforementioned collaborations that showed tension with the Standard Model expectation and could hint towards the existence of BSM physics [182–190].

⁴The standard full Run 2 data-set is larger than the 2018 B-parking data-set, hence is more constrained by the computational resources required for the event reconstruction.

4.4 Hadrons and Non-isolated Photons

After identifying muons, electrons and isolated photons, the remaining elements in the PF blocks are used to identify charged and neutral hadrons, as well as non-isolated photons (from eg. π^0 decays) [149].

Firstly, clusters from the ECAL and HCAL that are not linked to any KF track are interpreted as photons and neutral hadrons. For neutral hadrons inside jets it was observed that they deposit only 3% of the jet energy in the ECAL. ECAL clusters are therefore identified as photons, while HCAL clusters are identified as neutral hadrons. Outside the tracker acceptance ($|\eta| > 2.5$), this interpretation is not justified, since no distinction can be made between neutral hadrons and charged hadrons. In this region, ECAL-only clusters are therefore interpreted as photons, while ECAL clusters linked to an HCAL cluster are interpreted as hadrons.

The remaining objects are ECAL and HCAL clusters linked to tracks, and correspond to charged hadrons and possibly additional neutral candidates. The particles are identified by comparing the calibrated calorimetric energy to the sum of track momenta. In particular, if the calibrated calorimetric energy is larger than the sum of track momenta, then the excess is attributed to photons and neutral hadrons. In this case the tracks each give rise to a charged hadron with energy and momentum taken from the track, while the exact neutral particle content is determined from the size of the energy excess and recalibrations. If the calibrated calorimetric energy is similar to the sum of track momenta, then only charged hadrons are identified, with their momenta computed by including both Tracker and Calorimeter information in a χ^2 fit.

4.5 Higher-level Objects

After PF has identified all muons, e/γ objects and hadrons, they may be used to reconstruct higher level objects, such as (flavour-tagged) jets, hadronic taus and energy sums. The reconstruction and identification of these higher level objects relies on dedicated methods and is sensitive to contributions from pileup. The main methods used in CMS for the reconstruction and pileup mitigation are described in Sections 4.5.1 – 4.5.3. An alternative pileup mitigation method is described in Section 4.6.

4.5.1 Jets and Flavour Tagging

Jets are reconstructed by clustering PF candidates using the anti- k_T algorithm [140, 163]. In this algorithm two distances are defined between objects i, j and the beam B:

$$d_{ij} = \min(p_{Ti}^{-2}, p_{Tj}^{-2}) \frac{\Delta_{ij}^2}{R^2} \quad (4.1)$$

$$d_{iB} = p_{Ti}^{-2} \quad (4.2)$$

Here, $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, and k_{Ti} , y_i and ϕ_i represent the transverse momentum, rapidity and azimuth of object i , respectively. R is a distance parameter governing the size of the jet. The algorithm identifies the smallest distance between any two entities. If it corresponds to d_{ij} , then objects i and j are combined into a new object with coordinates computed as the energy-weighted average. Alternatively, if it corresponds to d_{iB} , then object i is promoted to a jet and the constituents removed from the list of entities. This procedure is repeated until no entities remain. The standard distance parameter in CMS is R=0.4 (AK4 jets), but also R=0.8 (AK8 jets) is commonly used to reconstruct Lorentz-boosted particles such as W, Z or Higgs bosons, where the decay products are collimated into a single *fat jet*.

Ideally the clustering algorithm only uses PF candidates that originate from the hard-scatter primary vertex, not those from pileup interactions. This is achieved only partially by ignoring PF candidates whose charged particle track is unambiguously associated to a pileup vertex (see Section 4.1.1). This procedure is known as charged hadron subtraction (CHS) [149, 164]. However, PF candidates whose track is not associated to any vertex are kept. Furthermore, charged candidates beyond the Tracker acceptance as well as neutral candidates cannot be associated to a primary vertex, hence are inadvertently included in the clustering procedure too.

A series of jet energy corrections (JEC) is applied to reconstructed jets in simulation and in data, in order to compensate for the remaining pileup contributions and account for detector noise, non-uniform detector response and differences between data and simulation [191, 192]. The jet energy scale (JES) is calibrated using a series of corrections as shown in Figure 4.7. The first layer (L1) of corrections after jet reconstruction is an *offset correction*, which aims to subtract the energy from electronics noise and residual pileup after CHS. This correction is parameterized as function of the (uncorrected) jet p_T , η , area and offset energy density,

and derived by looking at energy deposited in random cones (RC) in neutrino gun simulations and zero-bias data. Next, (p_T, η) -dependent correction factors (L2L3) are applied to improve the jet energy response. They are derived with QCD multi-jet simulations comparing the reconstructed and particle level jet p_T . Afterwards, only minor residual corrections are needed for jets in data. This (L2L3 residual) correction is extracted from $\gamma/Z + \text{jet}$ and multi-jet events, where one or more jets recoil against a reference object and the transverse momentum imbalance can be exploited to improve the jet energy response. An optional correction may furthermore be applied to improve the jet response for different flavours of the mother parton (ie. the parton that gave rise to the jet), accounting for varying jet fragmentation and particle composition. The effect of the various JES corrections is shown in Figure 4.8. After correcting the JES, also the jet energy resolution (JER) may be *smeared* in simulation to match that of data as described in Refs. [191, 192].

Jet flavour identification is a crucial capability for efficient selection (or rejection) of event topologies with hadronic activity, especially in the LHC environment. It is a multi-class classification problem that is generally approached by exploiting multiple variables sensitive to the jet fragmentation and particle composition. Also the lifetime of the mother particle is an important ingredient; Hadrons containing a b (c) quark have a typical lifetime of the order of 1.5 (1) ps, which means that they travel roughly 450 (300) μm before decaying [155]. As a result, the daughter tracks will have a relatively high impact parameter, defined as the distance of closest approach between the track and the primary vertex. The excellent transverse impact parameter resolution of CMS, driven by the innermost pixel layer, was found to be as low as 20 μm in 2017 [193], and allows the reconstruction of the displaced secondary vertex. As of Run 2, the inclusive vertex finding (IVF) algorithm is used to reconstruct secondary vertices corresponding to heavy-flavour decays, as described in Ref. [168]. Information from these secondary vertices is used in combination with track information and the presence of soft non-isolated

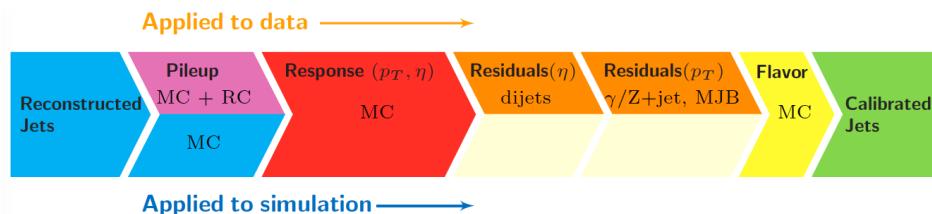


Figure 4.7: The factorized JES corrections applied to jets reconstructed in data (top half) and simulation (bottom half) [191].

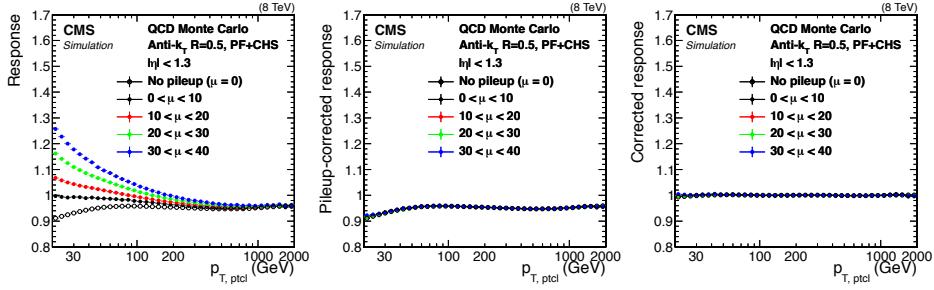


Figure 4.8: Jet p_T response as function of particle-level jet p_T , measured in QCD simulations, at various stages of the JES calibration: before any corrections (left), including only the offset corrections (middle) and after all JES corrections (right) [191].

leptons⁵ for multivariate identification of the jet flavour. Several taggers exist for b tagging; A version of the combined secondary vertex (CSV) tagger used in Run 1 was developed, that is based on deep machine learning (DeepCSV) [194, 195]. Alternatively, the DeepJet tagger was recently developed [196], that relies on less stringent selection of the jet constituents and a different model architecture, allowing it to exploit the full jet information and achieve higher performance, as is shown in Figure 4.9 (left). Given that both the DeepCSV and DeepJet tagger are multi-class classifiers, they are also used for c jet identification. However, the decreased discrimination power for c jets (versus light jets) results in generally worse performance than for b jets, as shown in Figure 4.9 (right).

4.5.2 Hadronic Taus

The tau is the heaviest lepton and the only type that can decay to hadrons. The branching ratios (BR) of the main decay modes are shown in Table 4.1. Leptonic decay modes make up about 35% of the total BR and result in a muon or electron (reconstructed as described in Sections 4.2 and 4.3) accompanied with two neutrinos that contribute to the missing transverse momentum. However, with roughly 65% BR the hadronic decay modes are dominant. Hadronic taus (τ_h) typically decay to either one charged hadron (referred to as 1-prong), or to three charged hadrons (3-prong), accompanied with a neutrino and up to two neutral pions.

The reconstruction of τ_h is performed with the *hadrons-plus-strips* (HPS) al-

⁵Roughly 20% (10%) of b (c) jets contain a low-energy electron or muon. Despite this low fraction, the presence of such leptons is used by some taggers to select a pure sample of heavy-flavour jets [168].

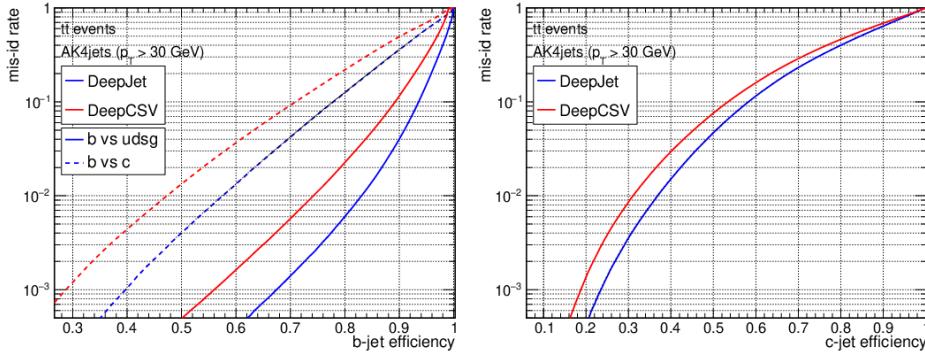


Figure 4.9: Performance of the DeepCSV and DeepJet taggers in fully hadronic $t\bar{t}$ events [196]. Left: performance for b vs. c classification (dashed lines) and b vs. light classification (solid lines). Right: performance for c vs. light classification.

gorithm [197–199]. It uses as input the constituents of reconstructed PF jets, requiring that the charged particles satisfy $p_T > 0.5$ GeV and transverse impact parameter $d_{xy} < 0.1$ cm, in order to ensure sufficient quality of the tracks and account for the notable lifetime of a genuine τ ($c\tau \simeq 87\mu\text{m}$). Next, a dynamic strip reconstruction aims to cluster all e/γ objects resulting from $\pi^0 \rightarrow \gamma\gamma$ and subsequent $\gamma \rightarrow e^+e^-$ decays associated to the original τ_h . This starts by defining a strip in the (η, ϕ) plane, centered around the e/γ object with highest p_T . Other e/γ objects are merged into the strip by iteratively adding the e/γ object with next-highest p_T that is located within the strip. The size of the strip is dynamically adjusted depending on the p_T of e/γ objects already part of the strip and the candidate object to be merged, up to $\Delta\eta \times \Delta\phi = 0.15 \times 0.3$. After merging a candidate, the center location of the strip is recalculated as the p_T weighted average

Decay mode	Resonance	BR (%)
$\tau^\pm \rightarrow e^\pm \nu_e \nu_\tau$		17.8
$\tau^\pm \rightarrow \mu^\pm \nu_\mu \nu_\tau$		17.4
$\tau^\pm \rightarrow h^\pm \nu_\tau$		11.5
$\tau^\pm \rightarrow h^\pm \pi^0 \nu_\tau$	$\rho(770)$	25.9
$\tau^\pm \rightarrow h^\pm \pi^0 \pi^0 \nu_\tau$	$a_1(1260)$	9.5
$\tau^\pm \rightarrow h^\pm h^\mp h^\pm \nu_\tau$	$a_1(1260)$	9.8
$\tau^\pm \rightarrow h^\pm h^\mp h^\pm \pi^0 \nu_\tau$		4.8
Other hadronic modes		3.3

Table 4.1: Decay modes of taus and their branching ratios [197]. Decays via intermediate resonances are indicated where appropriate.

Object and Event Reconstruction in CMS

of all e/γ constituents. This procedure continues until no e/γ candidates can be merged. If applicable, the e/γ object with highest p_T that is not associated to any strip seeds the reconstruction of a new strip, until all e/γ objects within the PF jet are used. The strips are then combined with the charged hadrons to reconstruct τ_h candidates for the five main hadronic decay modes of Table 4.1. The τ_h must have a charge of ± 1 and its constituents must lie within a p_T -dependent cone. If applicable, the constituents must also have an invariant mass compatible with the resonance associated to the τ_h decay mode hypothesis.

The identification of τ_h aims to distinguish genuine τ_h from jets, electrons or muons, and different discriminators exist to minimize each of the misidentification probabilities [197, 199]. For discrimination against jets, a multivariate discriminator based on BDTs relies on isolation information, variables sensitive to the τ_h lifetime, and e/γ object multiplicities. Also to distinguish τ_h from electrons dedicated BDTs were trained, that instead rely on the distribution of energy deposits in the ECAL, variables sensitive to the amount of bremsstrahlung, and object multiplicities. The $\tau_h \rightarrow \mu$ misidentification probability can be reduced by vetoing τ_h that align with signals in the muon chambers. More recently another classifier was developed based on deep machine learning (DeepTau) [200], that relies on a more inclusive object description to achieve superior performance, as is shown in Figure 4.10.

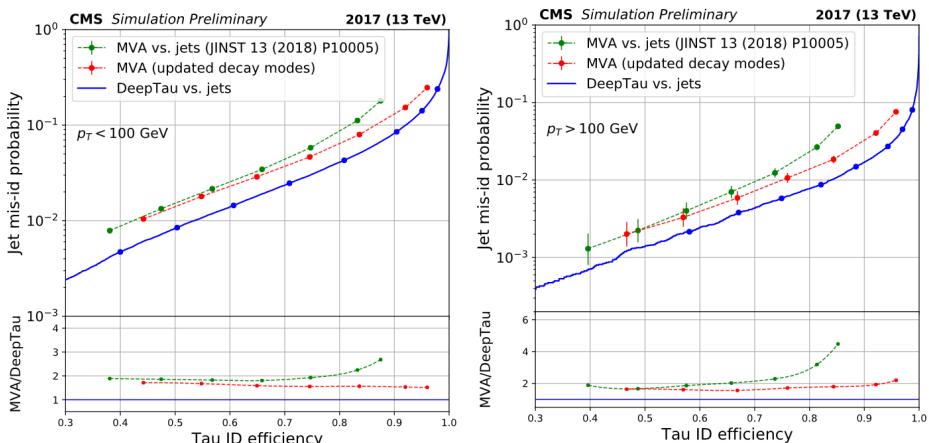


Figure 4.10: Performance of the main classifiers discriminating genuine τ_h from jets in simulated W+jets events, for $p_T < 100$ GeV (left) and $p_T > 100$ GeV (right) [200]. The green line corresponds to the BDT-based classifier of Ref. [197], which does not account for the $\tau_h \rightarrow h^\pm h^\mp h^\pm \pi^0$ decay mode. This mode was then added via a more inclusive decay mode definition used for both the red (BDT) and blue (DeepTau) lines.

4.5.3 Energy Sums

Energy sums are typically used both in offline analyses and – in simplified form – at trigger level to assess the scale and momentum imbalance of the hard interaction. The most common quantities are the scalar sum of the p_T of jets (H_T), as well as the missing transverse momentum (\vec{p}_T^{miss} , with magnitude p_T^{miss}). The \vec{p}_T^{miss} is computed as the negative vector sum of the \vec{p}_T of all visible final-state particles, imposing momentum conservation ($\vec{p}_T^{\text{in}} = \vec{p}_T^{\text{out}} = 0$). In particular the p_T^{miss} is a key quantity for many BSM searches, since true p_T^{miss} can only be induced by neutrinos and by stable and weakly interacting new physics states (such as Dark Matter candidates). However, both quantities are highly sensitive to the detector response, noise and the number of pileup interactions.

To correct for the above factors, the H_T can be computed simply by using the p_T of CHS PF jets after the JEC are applied (see Section 4.5.1), denoted as p_T^{JEC} , such that:

$$H_T = \sum_{j \in \text{jets}} p_{T,j}^{\text{JEC}} \quad (4.3)$$

Several corrections exist for the \vec{p}_T^{miss} . Firstly, it should be noted that the uncorrected (raw) \vec{p}_T^{miss} is the result of summing all final-state particles, being a mixture of particles from the hard-scatter and both charged and neutral pile-up contributions:

$$\vec{p}_T^{\text{miss, raw}} = - \sum_{i \in \text{PF objects}} \vec{p}_{T,i} \quad (4.4)$$

$$= - \sum_{i \in \text{HS}} \vec{p}_{T,i} - \sum_{i \in \text{neuPU}} \vec{p}_{T,i} - \sum_{i \in \text{chPU}} \vec{p}_{T,i} \quad (4.5)$$

The *Type-0 correction* serves solely to mitigate the pile-up terms. The charged contribution can be removed via CHS, which is assumed to be exact. The neutral component is taken to cancel the charged components at truth level, but its size is systematically measured off by a factor R^0 . The Type-0 corrected \vec{p}_T^{miss} , approximating the \vec{p}_T^{miss} associated to the hard-scatter only, can therefore be written as:

$$\vec{p}_T^{\text{miss, Type 0}} = \vec{p}_T^{\text{miss, raw}} + (R^0 - 1) \sum_{i \in \text{chPU}} \vec{p}_{T,i} \quad (4.6)$$

The more popular *Type-1 correction* consists of propagating the JEC of all clustered jets to the \vec{p}_T^{miss} . This is done by replacing the sum of \vec{p}_T of clustered particles with the sum of \vec{p}_T of calibrated jets:

$$\vec{p}_T^{\text{miss, Type 1}} = \vec{p}_T^{\text{miss, raw}} + \sum_{j \in \text{jets}} (\vec{p}_{T,j} - \vec{p}_{T,j}^{\text{JEC}}) \quad (4.7)$$

In contrast to the Type-0 corrections, the Type-1 corrections account also for the detector response and noise, but it does not correct contributions from unclustered particles. Although this can be solved by applying both Type-0 and Type-1 corrections simultaneously, this would imply that the pile-up components that are also clustered in jets, are over-corrected. Therefore, only Type-1 corrections are generally applied, and uncertainties may be assigned to account for the unclustered particles by varying their respective energy scales.

Finally, the *xy-Shift corrections* may be applied to account for the ϕ modulation of the \vec{p}_T^{miss} resulting from an-isotropic detector response, detector misalignment and displacement of the beam spot. It consists of shifting the origin in the transverse plane ($\vec{p}_T \rightarrow \vec{p}_T - \vec{c}$) for all n particles, such that:

$$\vec{p}_T^{\text{miss, xy}} = \vec{p}_T^{\text{miss, raw}} + n\vec{c} \quad (4.8)$$

The p_T^{miss} response and resolution is typically obtained from the momentum

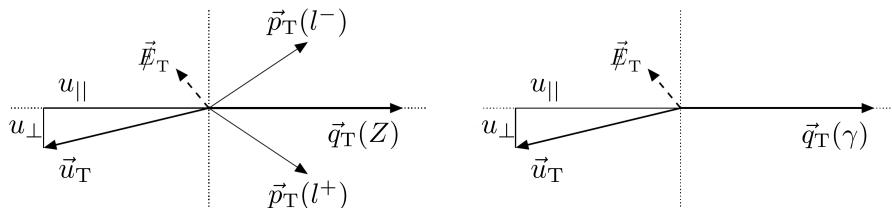


Figure 4.11: Kinematic configuration in the transverse plane for events with a $Z \rightarrow l^+l^-$ (left) or γ (right) recoiling against hadronic activity (denoted by \vec{u}_T), used to measure the p_T^{miss} response and resolution.

imbalance in events with a well-measured Z/γ boson recoiling against hadronic activity [201, 202]. In Figure 4.11 the \vec{p}_T of the Z/γ (hadronic recoil) is represented by \vec{q}_T (\vec{u}_T) such that momentum conservation implies $\vec{q}_T + \vec{u}_T + \vec{p}_T^{\text{miss}} = 0$. Since this topology provides a unique event axis with precise momentum scale due to the excellent reconstruction of leptons and photons, the response and resolution of the recoil may serve as a direct proxy for that of the p_T^{miss} . As shown in the figure, the parallel (perpendicular) component of \vec{u}_T projected onto the axis of \vec{q}_T is denoted as $u_{||}$ (u_{\perp}). The p_T^{miss} response is therefore $-\langle u_{||} \rangle / \langle q_T \rangle$, and the resolution is assessed via the root-mean-square of the $u_{||} + q_T$ and u_{\perp} distributions. The response and resolution as function of q_T is shown in Figure 4.12.

4.6 Pileup Per Particle Identification

Jets and energy sums are highly sensitive to the number of pile-up interactions. As mentioned in Section 4.5, several strategies may be employed to mitigate the pile-up contributions, such as the CHS and random-cone methods, or assumptions based on the true or reconstructed charged and neutral pile-up components. However, an alternative method was developed to further minimize the pile-up dependency of higher-level objects, referred to as *pile-up per particle identification* (PUPPI). [164, 203]. The PUPPI method relies on local particle distributions and tracking information to assign a weight to each PF candidate, reflecting the probability that it originates from the leading primary vertex.

Firstly, charged particles are assigned a weight based on tracking information. A weight of 1 is assigned to a charged PF candidate if it is associated to the hard-scatter vertex (ie. the track was used in the hard-scatter vertex fit), or if it is not associated to any vertex and the distance of closest approach between the track and the hard-scatter vertex in the z-direction is below 0.3 cm. In all other cases the charged PF candidate is assigned a weight of 0.

For neutral candidates, which cannot be associated to vertices due to the absence of tracks, the weight is assigned using a local particle distribution variable α . This variable is based on the notion that parton showers produced via QCD mechanisms exhibit a collinear structure, while pile-up interactions result in soft diffuse radiation that is distributed more homogeneously. To contrast the two mechanisms, α is defined for a given particle i as:

$$\alpha_i = \log \sum_{j \neq i, \Delta R_{ij} < 0.4} \left(\frac{p_{T,j}}{\Delta R_{ij}} \right)^2 \quad (4.9)$$

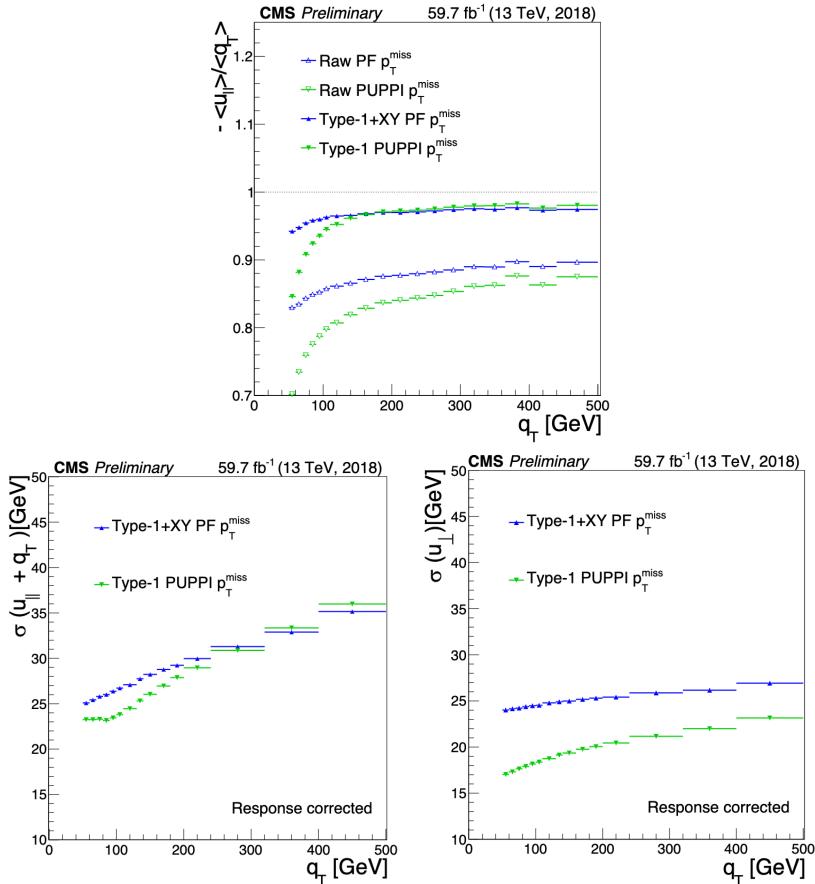


Figure 4.12: Response of p_T^{miss} (top) and resolution of the parallel (bottom-left) and perpendicular (bottom-right) components of the recoil, measured in γ +jets events [202]. PF p_T^{miss} (blue) is based on the standard PF candidates after CHS, while PUPPI p_T^{miss} (green) is based on PUPPI-weighted PF candidates as described in Section 4.6. In the 2018 data-taking conditions the PUPPI p_T^{miss} shows superior performance due to better pile-up mitigation of the PUPPI algorithm.

For a particle i within the Tracker acceptance ($|\eta_i| < 2.5$), the summation runs only over charged particles j associated to the hard-scatter vertex, while otherwise all charged and neutral candidates (within $\Delta R_{ij} < 0.4$) are included. The variable α is translated to a probability via a χ^2 approximation, using the expected α distribution of PU as reference:

$$\chi_i^2 = \frac{(\alpha_i - \bar{\alpha}_{\text{PU}})^2}{(\alpha_{\text{PU}}^{\text{RMS}})^2} \quad (4.10)$$

Here $\bar{\alpha}_{\text{PU}}$ and $\alpha_{\text{PU}}^{\text{RMS}}$ are the mean value and root-mean-square of α_i , respectively, computed using charged particles associated to pile-up vertices. Outside the Tracker acceptance, these values are computed by extrapolating the values for $|\eta| < 2.5$ using transfer factors derived from simulation, that account for varying detector granularity. Finally, the weight is obtained as:

$$w_i = F_{\chi^2, \text{NDF}=1}(\chi_i^2) \quad (4.11)$$

where $F_{\chi^2, \text{NDF}=1}$ is the cumulative distribution function, approximating the χ^2 distribution with one degree of freedom of all particles in the event.

Once all PF candidates are assigned a PUPPI weight, higher level objects can be computed based on the weighted four-vectors of the candidates. Objects with $w_i < 0.01$ are rejected to remove high-energy noise deposits and neutral candidates must pass $w_i p_{T,i} > (A + BN_{vtx})$. The latter condition reduces the remaining dependence of jet energies on the number of vertices (N_{vtx}), and A and B are tuned to optimize jet p_T and p_T^{miss} resolutions.

PUPPI-based jets are reconstructed by clustering PUPPI-weighted PF candidates using the anti- k_T clustering procedure described in Section 4.5.1. Due to the use of PUPPI to mitigate pile-up contributions, both CHS and the L1 offset correction to the JES are not needed. *PUPPI-based p_T^{miss}* is reconstructed via Equation 4.4, where the summation runs over all PUPPI-weighted PF candidates, and can be Type-1 corrected via Equation 4.7 using PUPPI-based jet calibrations. A comparison of the response and resolution for PF p_T^{miss} and PUPPI p_T^{miss} is shown in Figure 4.12.

Chapter 5

Search for Mass-compressed Prompt and Long-lived Electroweakinos with Soft Lepton Final States

As explained in Section 1.5 and Chapter 2, there have been numerous efforts at the CERN LHC to search for new physics phenomena that could answer to the shortcomings of the Standard Model, such as the nature of Dark Matter and the hierarchy problem. R-parity conserving Supersymmetry with sparticle masses at the electroweak energy scale provides an attractive solution, but no experimental evidence has been found so far. Nevertheless, there are still corners of the model parameter space that are within reach of the LHC physics potential, while being largely unprobed due to experimental challenges, and these *must* be explored.

This chapter reports a search for Supersymmetry with an unconventional signature that is particularly challenging to probe experimentally, involving both mass-compressed and long-lived sparticles. The targeted final state consists of low- p_T (soft) prompt or displaced leptons in addition to p_T^{miss} . The author was the main analyzer in this search, acting also as contact person on behalf of the analysis team within the CMS Collaboration.

The chapter is structured as follows: the signal models will be briefly summarized in Section 5.1, followed by a contextual note in Section 5.2 on the relevant preceding research in CMS. Section 5.4 describes the data and simulated samples, and Sections 5.5 and 5.6 show the selection criteria for the physics objects and events, respectively. The background estimations and treatment of systematic uncertainties are described in Sections 5.7 and 5.8. Finally, Section 5.9 shows the final results and interpretations in terms of the targeted signal models.

5.1 Signal Models

The potential Supersymmetry realizations targeted in this search have already been introduced in Section 2.4. Supersymmetry with compressed mass-spectra – where the mass-splitting Δm between sparticles is $O(10)$ GeV or less – constitutes a set of realizations that live in the corners of the model parameter space and are particularly challenging to probe. In such scenarios, when producing a sparticle, it may go through a cascade decay into the lightest sparticle that is assumed to be stable in R-parity conserving theories. Since the stable sparticle is heavier than the other SM decay products, it typically carries away most of the energy and momentum, while the remaining visible particles have only low momentum (referred to as “soft”). Consequently, final states consist of only soft particles and moderate p_T^{miss} , which do not pass standard trigger selections. Furthermore, tight kinematic cuts are required to suppress SM backgrounds such as QCD multi-jet production, which are typically several orders of magnitude larger than the signal expected from new physics. Due to the small signal acceptance, large statistical uncertainties and relatively low production cross-sections, such signal topologies remain largely unprobed. As such, Supersymmetry with compressed mass-spectra could have been easily missed and may actually still reside at the electroweak scale, within the discovery potential of the LHC.

Compressed mass-spectra in the electroweak sector of Supersymmetry may be realized in several ways. Firstly, higgsinos are manifestly compressed, with masses preferred to be $O(100)$ GeV to avoid fine-tuning of MSSM parameters and preserve naturalness [100, 101]. Therefore, if the lightest charginos ($\tilde{\chi}_1^{\pm}$) and neutralinos ($\tilde{\chi}_2^0, \tilde{\chi}_1^0$) are higgsino-like, then their production at the LHC would give rise to compressed mass-spectra signatures. Alternatively, if the lightest sparticles

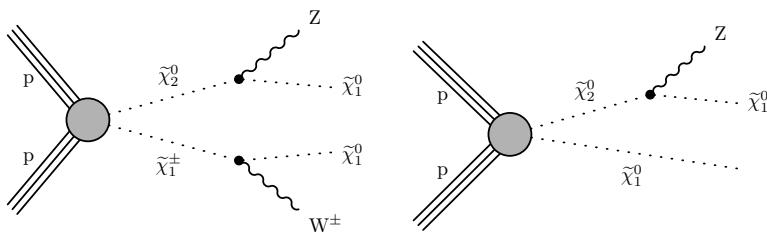


Figure 5.1: Targeted production and decay modes of electroweakinos in the wino-bino scenario (left) and the higgsino scenario (left and right).

are wino-like NLSP $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ with a bino-like LSP $\tilde{\chi}_1^0$, then compressed mass-spectra could be motivated by the wino-bino coannihilation process, which extends the description of the elementary particle annihilation and freeze-out in the early universe, predicting the observed amount of Dark Matter relic density [102, 103]. Both realizations are R-parity conserving and are henceforth referred to as the *higgsino scenario* and the *wino-bino scenario*.

In the higgsino scenario, the quasi-degenerate higgsino multiplet could lead to the production of any of the following combinations of sparticle pairs¹: $\tilde{\chi}_1^\pm\tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm\tilde{\chi}_2^0$, $\tilde{\chi}_1^\pm\tilde{\chi}_1^\pm$ or $\tilde{\chi}_1^0\tilde{\chi}_2^0$. In the wino-bino scenario, the dominant production modes are $\tilde{\chi}_1^\pm\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0\tilde{\chi}_1^\pm$, given the coupling structure of the electroweakinos to the SM gauge bosons [19]. In this search, the relevant signal topologies are $\tilde{\chi}_1^\pm\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0\tilde{\chi}_2^0$, as shown in Figure 5.1, with $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ both decaying a SM vector boson and a stable LSP $\tilde{\chi}_1^0$ assuming $\mathcal{B}(\tilde{\chi}_2^0 \rightarrow Z^*\tilde{\chi}_1^0) = \mathcal{B}(\tilde{\chi}_1^\pm \rightarrow W^*\tilde{\chi}_1^0) = 100\%$. Due to the small mass-splittings between the NLSP and LSP electroweakinos, the W and Z bosons from the NLSP decays are produced off-shell. This search targets in particular the leptonic decay of the off-shell Z^* boson, since this provides a crucial handle to reduce SM backgrounds. As such, the targeted final state consists of a pair of soft opposite-sign (OS) same-flavor (SF) leptons from the Z^* decay, with a possible third soft lepton from the W^* decay, in addition to p_T^{miss} from the pair of undetected $\tilde{\chi}_1^0$.

Both the above Supersymmetry scenarios have been previously targeted by preceding searches in CMS, as described in Section 5.2. These searches assume that the electroweakinos have negligible lifetime. However, considering wino-bino coannihilation is mini-split models – “mini-split” referring to the limited mass-difference between fermionic and bosonic superpartners to facilitate more easily the observed Higgs mass of 125 GeV – can lead to long-lived neutral wino states [108–110]. That is, in the topology of Figure 5.1 (left), the $\tilde{\chi}_2^0$ may become long-lived and its decay occur displaced with respect to its production vertex. The searches described in Section 5.2 have little to no sensitivity to such scenarios due to tight constraints on the provenance of the leptons with respect to the primary proton-proton (pp) interaction vertex. However, this scenario will be targeted in the search reported in the remainder of this chapter.

¹It can be shown that the $\tilde{\chi}_2^0\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0\tilde{\chi}_1^0$ production modes vanish [204], hence those are not considered.

5.2 Preceding Research in CMS

Within CMS, the leading search for mass-compressed electroweakinos is Ref. [89], which targets final states with two or three soft leptons and $p_T^{\text{miss}} > 125$ GeV using 137 fb^{-1} of pp collision data at $\sqrt{s} = 13$ TeV. Earlier analyses are reported in Ref. [205] (35.9 fb^{-1} at $\sqrt{s} = 13$ TeV) and Ref. [206] (19.7 fb^{-1} at $\sqrt{s} = 8$ TeV). Similar analyses were published by the ATLAS experiment [207–210].

The analysis strategy of Ref. [89] is summarized in the following, and lays the basis for the search for mass-compressed prompt *and* long-lived electroweakinos that is described in Section 5.3 onward.

5.2.1 Legacy Run 2 Analysis

Ref. [89] targets final states with two or three soft, light leptons (electrons or muons) and p_T^{miss} , making use of the full Run-2 data-set (137 fb^{-1} , collected during 2016–2018). To gain acceptance to signal scenarios with low mass-splittings between the NLSP and LSP, lower bounds on the electron (muon) p_T are set as low as 5 (3.5) GeV. The analysis strategy aims to select the leptons from the Z^* decay, hence the final state must include an opposite-sign same-flavor (OSSF) lepton pair with invariant mass $m_{\ell\ell}$ between 1 and 50 GeV. A dedicated search category is added in case a third soft lepton (from the W^* decay) is found.

The $m_{\ell\ell}$ of the reconstructed OSSF lepton pair is the key quantity in this search, since it serves as a proxy for the mass-splitting $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$. In Ref. [211] it is shown that the differential decay rate of $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z \rightarrow \ell\ell$ has a kinematic endpoint at $m_{\ell\ell} = \Delta m$ and depends on the structure of the couplings of the neutralinos to the Z boson. As such, the $m_{\ell\ell}$ distribution of signal is expected to have a peaking shape between $0 < m_{\ell\ell} < \Delta m$ GeV. This is in contrast to background events, noting that low-mass SM resonances (J/ψ and ν) are explicitly vetoed with mass-cuts. The $m_{\ell\ell}$ variable is therefore a great discriminator between signal and background. To improve the modeling of the dilepton mass-spectrum in signal, the electroweakino decays in MC simulated events – implemented using PYTHIA [212] – are reweighted according to the analytical description of Ref. [211]. Furthermore, as $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ (or $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$) decreases, some of the decay modes of the off-shell Z^* and W^* bosons become heavily suppressed. This differential effect for the 3-body decays of $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^\pm$ was calculated at tree-level with the SDECAY modules of the SUSYHIT 1.5A computational package [213] in order to account for the non-negligible second and third generation fermion masses. The differences with the PYTHIA simulated decays are shown in Figure 5.2 and are used

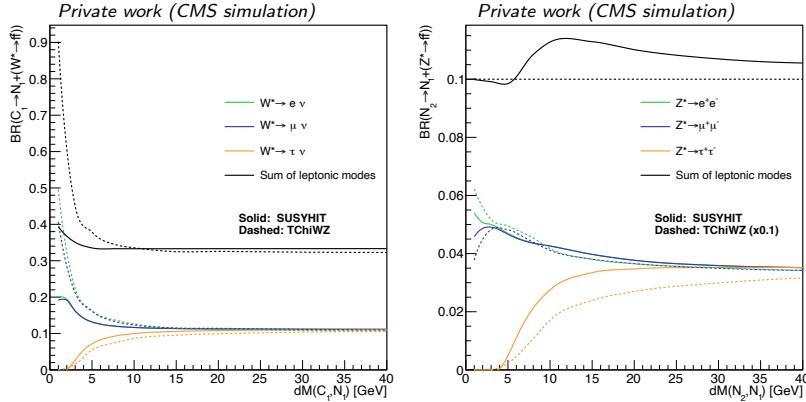


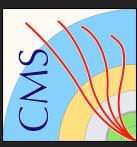
Figure 5.2: Left: Branching ratio of $\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 + W(\rightarrow \ell\nu)$ decays as function of $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$. Right: Branching ratio of $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 + Z(\rightarrow \ell\ell)$ decays as function of $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$. Both figures show the branching ratio modulation computed with SUSYHIT (solid lines) and the modulation in the MC simulated signal events obtained from PYTHIA (dashed lines).

to obtain correction factors that improve the modeling of the differential branching fractions in the simulated signal samples.

In the absence of high- p_T visible and invisible final state particles – the leptons are soft and the LSP $\tilde{\chi}_1^0$ is typically produced at rest, hence yields only moderate amounts of p_T^{miss} – the selection of signal events during data-taking is challenging, considering the overwhelming QCD background with a similar signature. Therefore, a jet from initial state radiation (ISR) is required in back-to-back configuration with the sparticle pair, in order to boost the sparticles and induce enough p_T^{miss} to pass the thresholds of p_T^{miss} -based triggers. In line with the trigger selection, the offline analysis requires events to satisfy $p_T^{\text{miss}} > 200$ GeV. To recover some of the unavoidable loss of signal acceptance from this strategy, a dedicated $\mu\mu + p_T^{\text{miss}}$ trigger was developed that allows the lower bound on offline p_T^{miss} to be decreased to 125 GeV with the additional requirement of two soft muons (offline $p_T > 5$ GeV).

Due to the selection of low- p_T leptons, the dominant background is from fake or non-prompt leptons, mainly from W +jets events. These are controlled via tight constraints on lepton parameters (such as the 3D impact parameter with respect to the primary vertex²) and estimated via data-driven methods. Other backgrounds include Drell-Yan (DY), $t\bar{t}$ and diboson production, each estimated with MC and scaled to data in dedicated control regions.

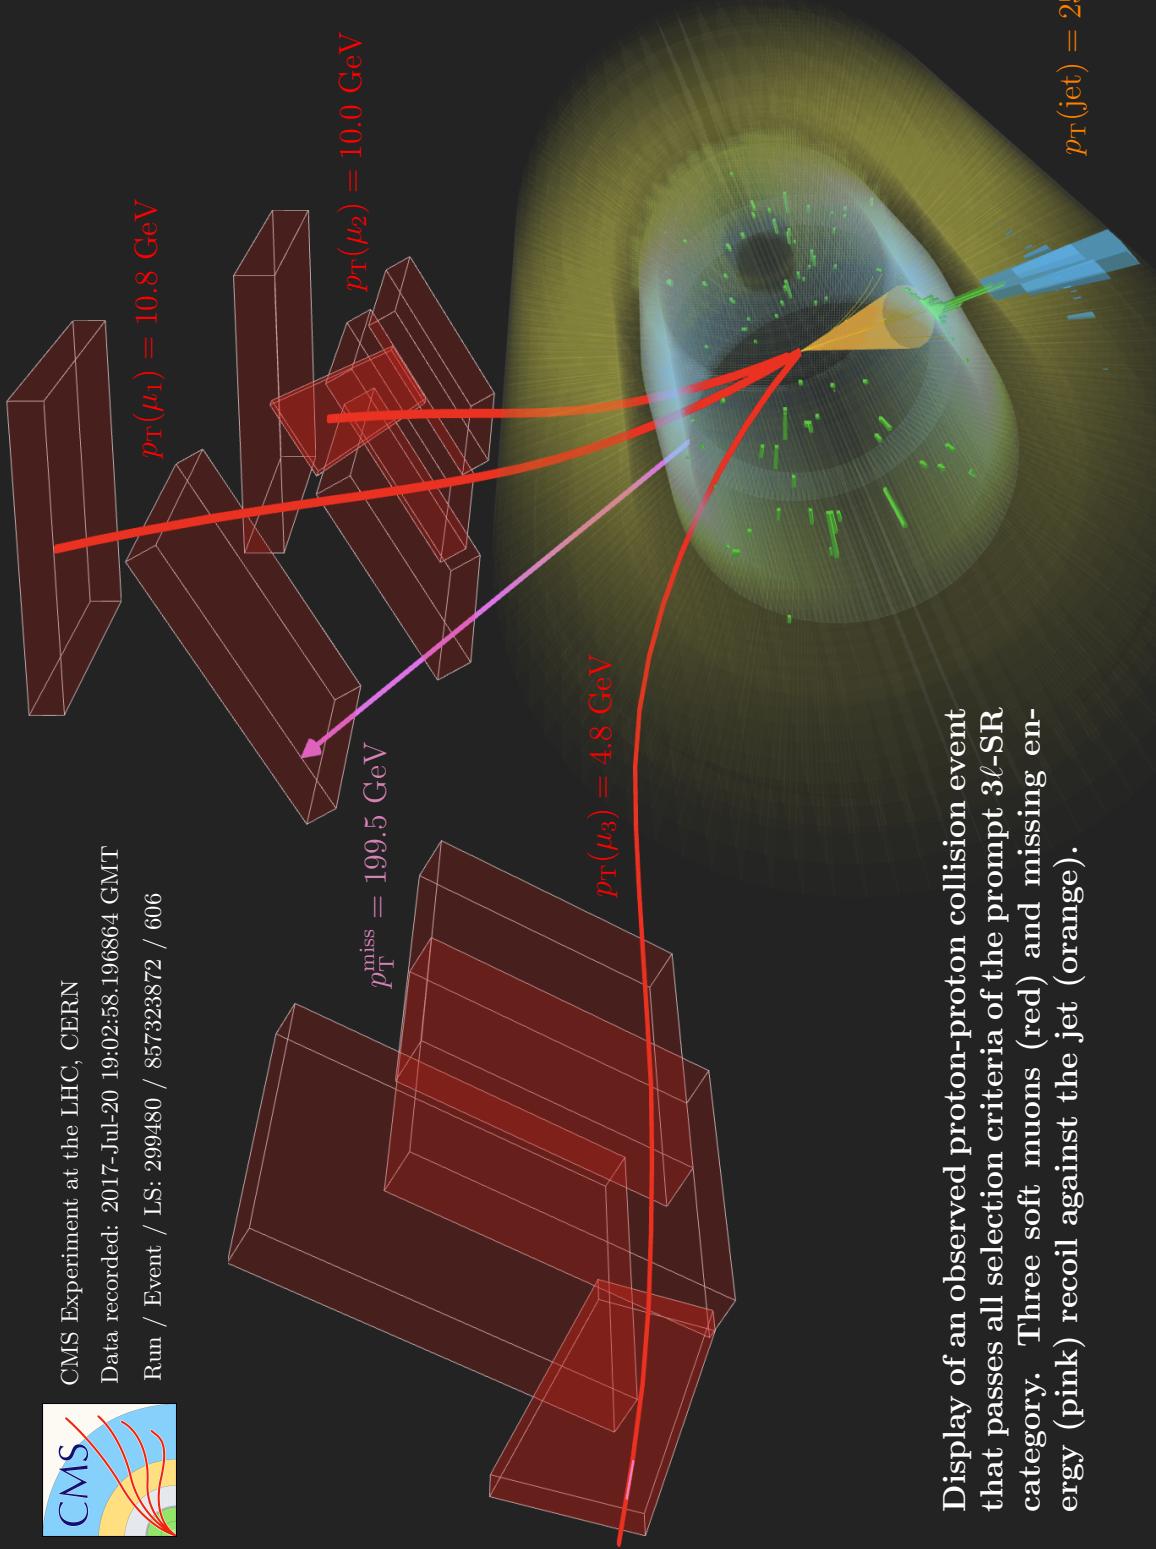
²This is the reason why this search has little sensitivity to long-lived sparticle decays.



CMS Experiment at the LHC, CERN

Data recorded: 2017-Jul-20 19:02:58.196864 GMT

Run / Event / LS: 299480 / 857323872 / 606



Display of an observed proton-proton collision event that passes all selection criteria of the prompt 3ℓ -SR category. Three soft muons (red) and missing energy (pink) recoil against the jet (orange).

$p_T(\text{jet}) = 253.6 \text{ GeV}$



CMS Experiment at the LHC, CERN

Data recorded: 2018-Jul-22 01:55:51.795904 GMT

Run / Event / LS: 320024 / 660213826 / 398

$$p_T^{\text{miss}} = 222.4 \text{ GeV}$$

$$p_T(\mu_1) = 7.0 \text{ GeV}$$

$$p_T(\mu_2) = 4.1 \text{ GeV}$$

$$p_T(\text{jet}) = 323.8 \text{ GeV}$$



Display of an observed proton-proton collision event that passes all selection criteria of the prompt 2μ -SR category. Two soft muons (red) are shown in addition to missing energy (pink) recoiling against the jet (orange).

The 2ℓ and 3ℓ search regions are divided into various p_T^{miss} bins and the corresponding $m_{\ell\ell}$ distributions, shown in Figure 5.3, are used in a maximum likelihood fit to data. No significant deviations between the background and observed data were observed, hence upper limits are set on the signal cross-section for the targeted models and a wide range of mass-hypotheses (Appendix A). At 95% confidence level, the search was able to exclude wino-like $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ masses up to 275 GeV at $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) = 10$ GeV and higgsino-like masses up to 205 GeV at $\Delta m = 7.5$ GeV. Sensitivity reached down to mass-splittings of roughly 3 GeV.

5.2.2 Parametric Signal Extraction

The signal extraction in Ref. [89] relies on a single $m_{\ell\ell}$ binning that is used uniformly across all search regions and for all signal mass-hypotheses. However, as mentioned above, the $m_{\ell\ell}$ signal shape depends strongly on the mass-splitting, particularly with the kinematic endpoint at $m_{\ell\ell} = \Delta m$. Consequently, the chosen $m_{\ell\ell}$ binning (that of Figure 5.3) may be sub-optimal for the considered mass-hypotheses. For example, for signal with $\Delta m = 3$ GeV, all events are expected to end up in the first $m_{\ell\ell}$ bin (ranging from 1 to 4 GeV), which means that the signal extraction profits less from the $m_{\ell\ell}$ shape, but rather on the total expected yields.

To better exploit the shape of the $m_{\ell\ell}$ distribution, an alternative signal extraction is developed that consists of optimizing the $m_{\ell\ell}$ binning for each Δm hypothesis, separately for each lepton category and p_T^{miss} bin. The optimization is an iterative procedure that starts with dividing the range $[1, \Delta m]$ GeV into four bins of equal signal quantiles and a fifth bin spanning $[\Delta m, 50]$ GeV, which is expected to be signal-depleted. The signal quantiles are based on the analytical expression of the $m_{\ell\ell}$ shape in Ref. [211]. The distribution of events per bin is therefore by construction uniform below Δm (modulo effects from resolution or statistical fluctuations), as shown in Figure 5.4.

The second step consists of checking that none of the resulting bins introduces large statistical fluctuations of the expected SM background, in order to avoid an artificial increase in the sensitivity. The figure of merit to evaluate the statistical power of estimating the SM background per $m_{\ell\ell}$ bin is the *number of equivalent unweighted events* N_{eq} , defined as:

$$N_{eq} = \frac{(\sum_{i \in \text{events}} \text{weight}_i)^2}{\sum_{i \in \text{events}} \text{weight}_i^2} = \frac{N_{\text{unweighted}} \langle \text{weight} \rangle^2}{\langle \text{weight}^2 \rangle} \quad (5.1)$$

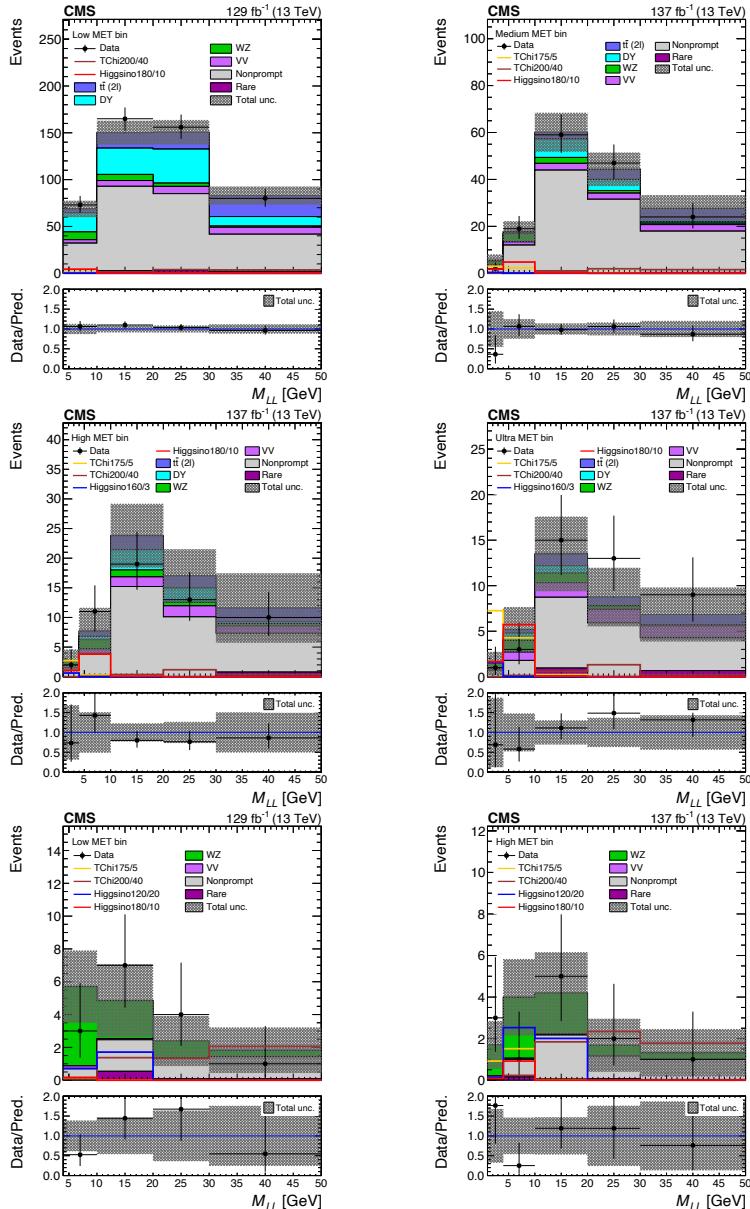


Figure 5.3: Distributions of the $m_{\ell\ell}$ variable for the 2ℓ (first and second row) and 3ℓ search regions [89]. The figures correspond to different MET categories (indicated in the legends), for which the selections are shown in Table 5.6.

where the sum runs over all unweighted background events and the weight_{*i*} includes all possible contributions to the weight of event *i* (such as luminosity scaling, trigger scale-factors, etc.). If a given $m_{\ell\ell}$ bin is populated by only a single background process, then N_{eq} tends to $N_{\text{unweighted}}$ (since $\langle \text{weight} \rangle^2 \approx \langle \text{weight}^2 \rangle$). In contrast, if it is populated by multiple processes, then N_{eq} tends to that of the most dominant background process. Therefore, setting a lower bound on N_{eq} prevents tuning the $m_{\ell\ell}$ binnings on fluctuations of the background, while taking into account both the relative importance of the various contributions and the statistical precision in their predictions. The lower bound itself was chosen to be $N_{eq} = 1.5$, based on the lowest value obtained with the static binnings in Figure 5.3. If for any of the bins in $[1, \Delta m]$ GeV the expected background has $N_{eq} < 1.5$, then the number of bins in that range is reduced by one, and the binning is rederived. This procedure is repeated until a binning is found that satisfies $N_{eq} > 1.5$ for all $m_{\ell\ell}$ bins (separately for each year), or results in the minimal two-bin case $[1, \Delta m, 50]$ GeV. The parameteric $m_{\ell\ell}$ binnings are derived for the SRs only, the CRs rely on the default static binning.

For each Δm hypothesis, six parametric $m_{\ell\ell}$ binnings have been derived (one for each SR). As an example, post-fit SR $m_{\ell\ell}$ distributions with parametric binnings derived for $\Delta m = 20$ GeV are shown in Figure 5.5. The shape of the background distributions may be compared to those in Figure 5.3 (the only difference being the $m_{\ell\ell}$ binning).

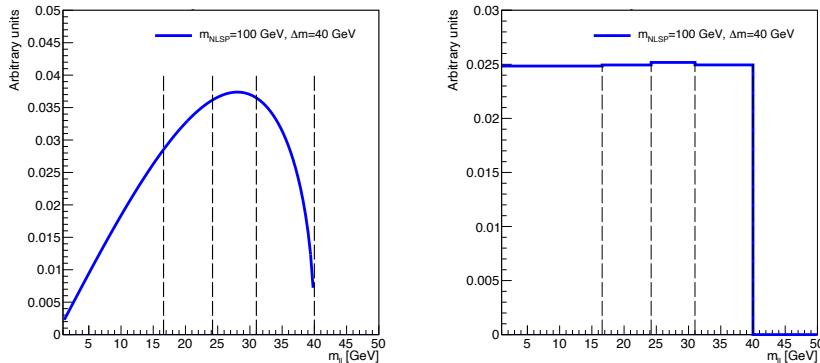


Figure 5.4: Left: Analytical $m_{\ell\ell}$ shape of signal with $m_{\tilde{\chi}_2^0} = 100$ GeV and $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) = 40$ GeV. The dashed lines indicate the $m_{\ell\ell}$ binning used in Ref. [89]. Right: the same signal shape, distributed according to the parametric $m_{\ell\ell}$ binning designed for $\Delta m = 40$ GeV, shown with the dashed lines. Note that the signal shape with the parametric binning is approximately uniform at $1 < m_{\ell\ell} < \Delta m$ GeV.

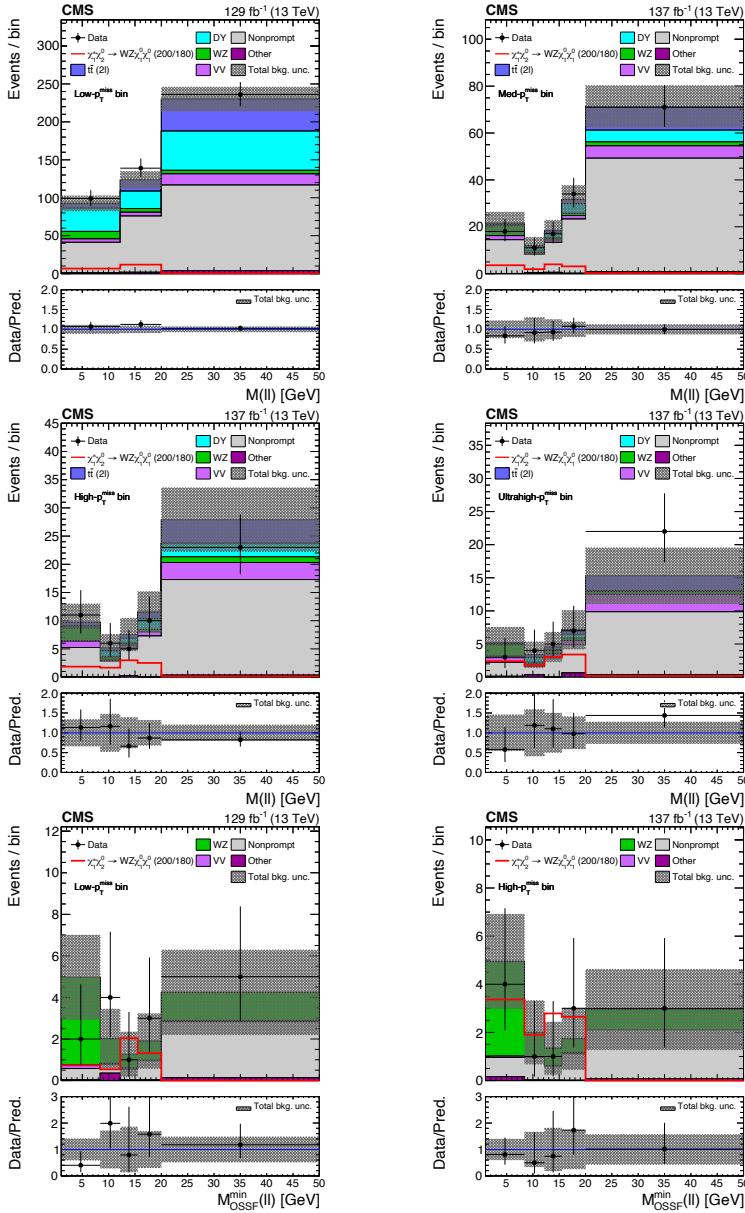


Figure 5.5: Distributions of the $m_{\ell\ell}$ variable for the 2ℓ (first and second row) and 3ℓ search regions, based on parametric $m_{\ell\ell}$ binnings designed for mass-splitting with $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) = 20 \text{ GeV}$ [214]. The figures correspond to different MET categories (indicated in the legends), for which the selections are shown in Table 5.6

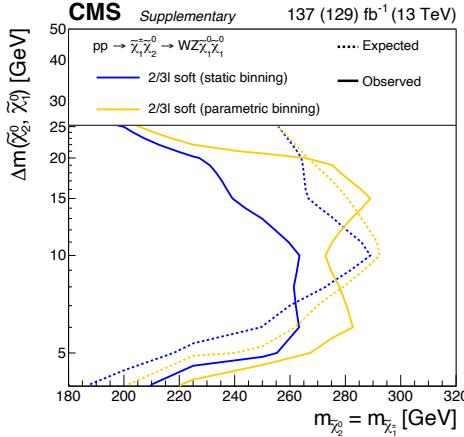


Figure 5.6: The 95% confidence level exclusion boundaries for wino-like chargino-neutralino production, as obtained with the *static* (unoptimized) $m_{\ell\ell}$ binnings of Ref. [89] and the *parametric* $m_{\ell\ell}$ binnings [214]. The parametric binnings increase the signal sensitivity over the full model parameter space considered in this search.

The gain of sensitivity with the parametric binnings by comparing the expected exclusion limits to those computed with the static $m_{\ell\ell}$ binnings, as shown in Figure 5.6. The expected sensitivity is improved over the full parameter space, with the largest gains at $\Delta m < 20$ GeV, where the expected $\tilde{\chi}_2^0$ mass-exclusion is increased by up to 25 GeV.

5.3 Overview of the New Analysis Strategy

As mentioned above, Ref. [89] had a peak sensitivity to signals with a mass-splitting of 10 GeV. However, the phase-space for $\Delta m < 5$ GeV was hardly accessible due to the lower p_T bounds for electrons (muons) of 5 (3.5) GeV, as shown in Figure 5.3 (left). Lowering the p_T thresholds to increase acceptance to smaller mass-splittings is very challenging; It implies significantly larger background contributions from hadronic activity, while also the reconstruction efficiency for both electrons and muons rapidly decreases. Furthermore, to control the already large hadronic background from non-prompt heavy flavor decays, Ref. [89] imposes tight requirements on the lepton impact parameters. Consequently, the analysis had little to no sensitivity to new physics scenarios with long-lived decays.

The search presented in the remainder of this chapter follows in the footsteps of Ref. [89], employing a similar analysis strategy, but expanding the reach of the analysis in two directions:

5.3 Overview of the New Analysis Strategy

- **More compressed new physics states** - To gain acceptance to signals with lower mass-splittings, the p_T thresholds *must* be decreased. With the conventional object reconstructions this is hardly possible. However, an alternative *LowPT electron* reconstruction allows lowering the p_T bounds for electrons down to 1 GeV, which opens up previously inaccessible phase-space where Supersymmetry may reside. As described in Section 4.3.5, this LowPT electron reconstruction was originally designed to maximize the $B \rightarrow X + ee$ data-set for eg. measurements of leptonic heavy flavor decay ratios. However, this chapter describes one of the first use-cases of these objects in a *search*.
- **Long-lived new physics states** - Figure 5.3 (right) shows the phase-space populated by wino-bino coannihilation scenarios in the mini-split model. In these scenarios the $\tilde{\chi}_2^0$ can become long-lived and decay to a pair of soft *and* displaced lepton pairs. To gain acceptance to these scenarios, a new search category is designed based on the presence of a soft displaced muon pair, which is constructed by refitting standard muon tracks to a common displaced vertex. Since mass-splittings of several GeV are required in order to be able to reconstruct the muons, the maximum $\tilde{\chi}_2^0$ life-time considered in this search is roughly 100 mm.

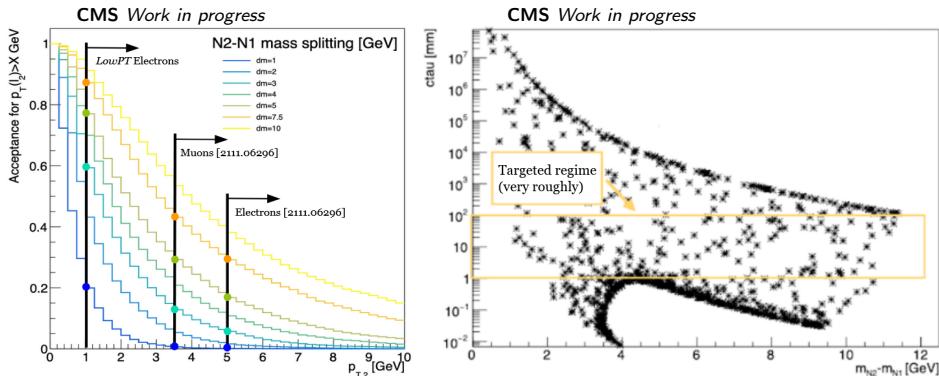


Figure 5.7: Left: Acceptance to compressed signal hypotheses as function of the p_T threshold on the sub-leading lepton. The black lines indicate the electron and muon thresholds imposed in Ref. [89], as well as the threshold for the alternative LowPT electron reconstruction. Right: Phase-space populated by wino-bino coannihilation scenarios in the mini-split Supersymmetry model, as calculated with SOFTSUSY 4.1.8 [215]. The points are based on a random scan where $M1=100 \text{ GeV}$, $M2$ in $(90,100) \text{ GeV}$, μ in $(500 \text{ GeV}, 25 \text{ TeV})$ and $\tan(\beta)=30$. The region of interest for the displaced search category is highlighted in yellow.

5.4 Data and Simulated Samples

Similar to Ref. [89], the analysis presented in the remainder of this chapter relies on the full Run 2 data-set as well, consisting of approximately 138 fb^{-1} of pp collision data collected at $\sqrt{s} = 13 \text{ TeV}$ by the CMS experiment. The total integrated luminosity per year amounts to 36.3 , 41.5 and 59.8 fb^{-1} for 2016, 2017 and 2018, respectively. The trigger strategy and performance are described in Section 5.4.1, while an overview of the MC simulated signal and background samples is given in Section 5.4.2. For the 2016 data-taking period, it was found that highly ionizing particles (HIP) could cause a temporary saturation of the APV25 front-end chips in the silicon strip Tracker sensors, leading to significant dead-time of the readout system. This issue was present during the first 19.5 fb^{-1} (labeled as “2016APV”), but mitigated during the remaining 16.8 fb^{-1} collected in 2016. The HIP(-mitigation) effect is also simulated in the MC samples.

5.4.1 Trigger Strategy

The trigger strategy in this analysis relies broadly on two types of triggers, both requiring a certain amount of p_T^{miss} in the event, as is common to searches for R-parity conserving realizations of Supersymmetry. In the Level-1 Trigger (L1T), p_T^{miss} is computed with calorimeter based physics objects only (Section 3.3.1), which therefore does not account for presence of muons. The p_T^{miss} calculated excluding the muons is denoted as $p_T^{\text{miss,corr}}$. This quantity is also used in offline analysis for compatibility with the trigger selection. The term “MET” is used throughout this chapter to denote either or both of these quantities. All triggers used in the analysis are summarized in Table 5.1.

The first trigger type is a pure $p_T^{\text{miss,corr}}$ based High-Level Trigger (HLT) path, used to probe regions with offline $p_T^{\text{miss,corr}} > 200 \text{ GeV}$ (referred to as “high-MET”). The second type is a complementary $\mu\mu + p_T^{\text{miss}}$ based trigger, designed [205] to accept events with offline $p_T^{\text{miss}} > 125 \text{ GeV}$ and $125 < p_T^{\text{miss,corr}} < 200 \text{ GeV}$ (referred to as “low-MET”), hence lowering the $p_T^{\text{miss,corr}}$ threshold with the additional requirement of two low- p_T muons. Finally, since the low-MET triggers impose an upper bound on the $m_{\ell\ell}$ of the muon pair, another set of triggers is used select events for the low-MET part of the *WZ enriched region*, constructed in the analysis to control the $\text{WZ}^{(*)}$ background (Section 5.7). These are the lowest-threshold unprescaled single lepton (electron or muon) triggers.

The performances of the low- and high-MET trigger paths are shown below, measured both in data and in MC simulated events (generally a mixture of Drell-

5.4 Data and Simulated Samples

Year	Region	HLT Path	Luminosity	HLT Cuts	Associated Offline Cuts
2016	High-MET SR, CR	HLT_PFMETNoMu120_PFMHTNoMu120_IDTight	36.3 fb^{-1}	$p_T^{\text{miss,corr}} > 120 \text{ GeV}$	$p_T^{\text{miss,corr}} > 200 \text{ GeV}$
	Low-MET SR & CR	HLT_DoubleMu3_PFMET50	33.5 fb^{-1}	$p_T^{\text{miss}} > 50 \text{ GeV}$ $m_{\ell\ell} < 60 \text{ GeV}$ $p_T > 3 \text{ GeV}$ $DCA < 0.5 \text{ cm}$	$p_T^{\text{miss}} > 125 \text{ GeV}$ $125 \text{ GeV} < p_T^{\text{miss,corr}} < 200 \text{ GeV}$ $4 < m_{\ell\ell} < 50 \text{ GeV}$ $p_T > 5 \text{ GeV}$ $DCA < 0.2 \text{ cm}^\dagger$
	Low-MET WZ-CR	HLT_IsoMu24	36.3 fb^{-1}	$p_T > 24 \text{ GeV}$	$p_T > 30 \text{ GeV}$
2017	High-MET SR, CR	HLT_PFMETNoMu120_PFMHTNoMu120_IDTight_PFH160	41.5 fb^{-1}	$p_T^{\text{miss,corr}} > 120 \text{ GeV}$	$p_T^{\text{miss,corr}} > 200 \text{ GeV}$
	Low-MET SR & CR	HLT_DoubleMu3_DZ_PFMET50_PFMHT60	36.7 fb^{-1}	$p_T^{\text{miss}} > 50 \text{ GeV}$ $3.8 < m_{\ell\ell} < 60 \text{ GeV}$ $p_T > 3 \text{ GeV}$ $d_z < 0.2 \text{ cm}$	$p_T^{\text{miss}} > 125 \text{ GeV}$ $125 \text{ GeV} < p_T^{\text{miss,corr}} < 200 \text{ GeV}$ $4 < m_{\ell\ell} < 50 \text{ GeV}$ $p_T > 5 \text{ GeV}$ $d_z < 0.1 \text{ cm}$
	Low-MET WZ-CR	HLT_IsoMu27	41.5 fb^{-1}	$p_T > 27 \text{ GeV}$	$p_T > 30 \text{ GeV}$
2018	High-MET SR, CR	HLT_Ele35_WPTight_Gsf	41.5 fb^{-1}	$p_T > 35 \text{ GeV}$	$p_T > 37 \text{ GeV}$
	Low-MET SR & CR	HLT_PFMETNoMu120_PFMHTNoMu120_IDTight_PFH160	59.8 fb^{-1}	$p_T^{\text{miss,corr}} > 120 \text{ GeV}$	$p_T^{\text{miss,corr}} > 200 \text{ GeV}$
	Low-MET WZ-CR	HLT_DoubleMu3_DCA_PFMET50_PFMHT60	59.3 fb^{-1}	$p_T^{\text{miss,corr}} > 50 \text{ GeV}$ $3.8 < m_{\ell\ell} < 60 \text{ GeV}$ $p_T > 3 \text{ GeV}$ $DCA < 0.5 \text{ cm}$	$p_T^{\text{miss}} > 125 \text{ GeV}$ $125 \text{ GeV} < p_T^{\text{miss,corr}} < 200 \text{ GeV}$ $4 < m_{\ell\ell} < 50 \text{ GeV}$ $p_T > 5 \text{ GeV}$ $DCA < 0.2 \text{ cm}^\dagger$

Table 5.1: List of HLT paths used during the 2016 – 2018 data-taking period to select events for the various analysis regions (which are defined in Sections 5.6 and 5.7). “CR” refers to all the control regions excluding the low-MET WZ-CR, which relies on different triggers. The two last columns show the relevant HLT cuts and associated offline cuts. The offline DCA cuts marked with † are applied in the displaced analysis only. For 2017 the DCA-version of the $\mu\mu + p_T^{\text{miss}}$ trigger recorded only 7.7 fb^{-1} . Therefore, the Δz -version is used in the 2017 prompt analysis, while no Low-MET region is present in the 2017 displaced analysis.

Yan, $t\bar{t}$ and diboson processes). The efficiencies are measured as the ratio of offline selected events to also pass the online trigger selection. Since the performance in simulated events might differ from that in data, correction factors are derived and applied as weights to simulated events.

Low-MET Trigger Paths

The low-MET $\mu\mu + p_T^{\text{miss}}$ based HLT triggers are cross-seeded by double- μ and $p_T^{\text{miss,corr}}$ paths at the L1T and a double- μ and p_T^{miss} path at the HLT. As indicated in Table 5.1, the HLT paths also feature extra requirements on the muon pair: Opposite charge, an upper cut on the dimuon mass (and lower bounds in 2017 and 2018), and an upper cut on the spatial separation of the muon tracks. Two versions of this trigger exist, based on different definitions of the latter, being either the *distance of closest approach* (DCA, corresponding to the smallest (3D) distance of the two muon tracks) or the Δz (corresponding to the distance of either muon track to the primary vertex in the z -direction). Since this analysis targets both prompt *and* displaced leptons, preference was given to the DCA version, which maintains efficiency also to displaced signatures. However, during 2017 the DCA-based trigger recorded only 7.7 fb^{-1} , which motivated the use of the Δz -based version instead (for 2017 only). Consequently, no low-MET search region was defined for the 2017 analysis (Section 5.6). Furthermore, since this $\mu\mu + p_T^{\text{miss}}$ based HLT path was not part of the trigger menu before the end of June 2016 technical stop, the collected luminosity amounts to $L_{\text{int}} = 33.5 \text{ fb}^{-1}$.

The efficiency is factorized into the various trigger selections as:

$$\epsilon = \epsilon_{\mu_1} \epsilon_{\mu_2} \times \epsilon_{\mu \text{ Distance}} \times \epsilon_{\text{MET}} \quad (5.2)$$

where it should be noted that the requirement on the dimuon invariant mass is not included, since its inefficiency was found to be negligible for $4 < m_{\ell\ell} < 56 \text{ GeV}$ (Figure 5.8). The efficiencies for each component in data and MC simulated events are shown for each period of data-taking in Figures 5.9-5.12.

The efficiencies of the MET part (ϵ_{MET}) are measured with events containing two muons that pass the tight offline object selections (defined in Section 5.5), firing a double muon trigger with online $p_T > 17(8) \text{ GeV}$ for $\mu_1(\mu_2)$. The muons must furthermore pass offline $10 < p_T < 50 \text{ GeV}$ and $4 < m_{\ell\ell} < 56 \text{ GeV}$, similar to the offline events selections of the analysis, making also the DCA (or Δz) component fully efficient. The efficiency ϵ_{MET} is then defined as the fraction of events also passing the $\mu\mu + p_T^{\text{miss}}$ trigger. The efficiency is parameterized as

function of the offline reconstructed p_T^{miss} and $p_T^{\text{miss,corr}}$, where it may be noted that the lower bounds for selected events (> 125 GeV for both quantities) are situated on the *turn-on* of the trigger efficiency.

The efficiencies for the leptonic part ($\epsilon_{\mu_{1,2}}$) are computed with the *tag-and-probe* method, using events that pass the L1T $p_T^{\text{miss,corr}}$ condition and contain two muons satisfying the tight selection criteria. One of the muons must be isolated and satisfy $p_T > 27(24)$ GeV for 2017 (2016 and 2018) at the HLT, and is hereby identified as the tag muon. The second (probe) muon is used to compute the efficiency as the fraction of probe muons passing the muon selections of the $\mu\mu + p_T^{\text{miss}}$ path. The efficiencies are parameterized as function of the offline p_T and η of the muons. Specifically for 2016, this term includes also the DCA efficiency. To include the DCA efficiency in Equation 5.2 only once, the equation is modified accordingly to:

$$\epsilon = \frac{\epsilon_{\mu_1} \epsilon_{\mu_2}}{\epsilon_{\text{DCA}}} \times \epsilon_{p_T^{\text{miss,corr}}} \quad (5.3)$$

For computing the DCA (2016) and Δz (2017) efficiency ($\epsilon_{\mu \text{ Distance}}$), events are required to pass a double muon trigger with $p_T > 17(8)$ GeV for $\mu_1(\mu_2)$ as well as the $\mu\mu + p_T^{\text{miss}}$ L1T seed. Both muons must pass the tight identification criteria of the analysis, have opposite charge and $4 < m_{\ell\ell} < 56$ GeV. The efficiency is then calculated as the fraction of events passing the DCA (or Δz) requirement, and is parametrized as function of the offline η of both muons. In the corresponding efficiency maps of Figures 5.9-5.12 the muons are chosen such that $\eta(\mu_1) > \eta(\mu_2)$, hence the upper left half of the maps are empty. For 2018, the DCA efficiency is derived from the efficiency measurements as function of displacement (described below). Due to the relative uniformity of the efficiency maps, a single value is chosen by averaging over the phase-space of the measurement. In data this amounts to 85.6%, 95.0% and 99.5% for 2016 APV, 2016 non-APV and 2017, respectively. In MC this amounts to 91.4%, 96.0% and 99.3% for 2016 APV, 2016 non-APV and 2017, respectively.

Finally, the efficiencies discussed above may be used for events with final state leptons that are closely matched to the primary vertex. However, the same triggers will also be used for the displaced analysis regions, which targets long-lived decays with leptons from a displaced vertex. This means that in particular the lepton p_T -, mass- and DCA-components have to be derived as function of displacement. For the MET-component of the low-MET trigger, as well as the high-MET trigger, the same scale-factors can be used as for the prompt analysis regions, since these efficiencies are not affected by the lepton displacement. To measure the efficiency

of the lepton p_T -components of the low-MET trigger as function of displacement, fits to the dimuon mass are performed in a region rich in J/ψ . For the MC measurement, simulated samples with high p_T $Z \rightarrow \nu\nu$ are used, while the data measurement relies on a mixture of events collected by orthogonal triggers based on p_T^{miss} , jets and e/γ objects. Events are required to pass the L1 MET and HT conditions and have one reconstructed muon pair. Both legs of the dimuon must pass the displaced muon identification criteria in addition to $\text{IP}_{3D} > 0.0175$ cm and $5 < p_T < 30$ GeV, as per the offline object selections (Section 5.5). The efficiency is then defined as the fraction of events that passes the filters requiring two L1 muons ($p_T > 3$ GeV) and ≥ 2 HLT muons ($p_T > 3$ GeV). The combined L1 and HLT efficiency of the lepton-component shows a flat dependency on the transverse displacement, significance of the transverse displacement, and leading muon p_T as shown in Figure 5.13 (top and middle rows). The data-to-MC scale factor for this measurement amounts to 0.98 ± 0.01 . For the mass- and DCA-components in 2018, a slightly different strategy was used. It should firstly be noted that the L1 mass filter (> 4 GeV) rejects nearly all J/ψ events. The efficiency of these components is therefore defined instead as the fraction of events passing the DCA (+mass) filter of a similar J/ψ -specific trigger. Events are selected with a mixture of single muon triggers and must contain two L1 and HLT muons ($p_T > 3$ GeV) and a muon pair with $2.9 < m_{\ell\ell} < 3.3$ GeV. MC events are taken from the same $Z \rightarrow \nu\nu$ sample, and no specific trigger path is required in the denominator to allow for sufficient statistics. Although this yields a softer p_T spectrum than in data, this is acceptable due to the weak p_T dependency of the DCA (+mass) efficiency. No fit is performed to constrain background due to the absence of side-bands resulting from the tight mass cuts. The HLT efficiency of the combined mass- and DCA-components are shown as function of the vertex transverse displacement and its significance in Figure 5.13 (bottom row). The uniform behavior of data and MC leads to a scale-factor of 1.00 ± 0.02 .

High-MET Trigger Paths

The high-MET trigger paths correspond to the lowest-threshold unprescaled version of each year of data-taking, which are seeded by pure $p_T^{\text{miss,corr}}$ paths at the L1T. Their efficiencies are calculated with events that are selected by triggers uncorrelated with $p_T^{\text{miss,corr}}$, specifically single muon paths with $p_T > 27(24)$ GeV for 2017 (2016 and 2018). The offline H_T is required to be higher than 100 GeV, in accordance with the analysis selection. The efficiencies are shown in Figure 5.14. To better estimate the efficiency at the turn-on and the plateau, the following

parametrization is used:

$$\epsilon = 0.5 \times \epsilon_\infty \times \left(\text{Erf} \left(\frac{p_T^{\text{miss,corr}} - \mu}{\sigma} \right) + 1 \right) \quad (5.4)$$

5.4.2 Monte-Carlo Simulations

Simulated events are used to predict event yields from the relevant SM background processes as well as the targeted signal hypotheses. Leading background components - $t\bar{t}$, Drell-Yan and $W+jets$ (the latter two in bins of H_T) - are simulated at leading-order (LO) in perturbative QCD with the MADGRAPH [216] event generator. Other samples (diboson, single top quark, and rare processes) are produced at next-to-leading order (NLO) with AMC@NLO [217] or POWHEG [218–222]. The NNPDF3.1 [223] LO and NLO parton distribution function sets are used, with accuracy matching that of the matrix element generator. Showering and hadronization is done by PYTHIA [212], which is also used for the underlying event description using the CP5 tune [224]. Finally, the simulation of the CMS detector is performed with the GEANT [225] package.

Signal samples for the wino-bino scenario with promptly decaying $\tilde{\chi}_2^0$, corresponding to the topology of Figure 5.1 (left), are generated using a similar setup; Pair production of the supersymmetric particles is modeled at LO with MADGRAPH, whereas their decay, parton showering and hadronization is simulated with PYTHIA. Detector simulation is performed with GEANT. The simulated signal hypotheses constitute a grid of mass-points that gets finer towards lower

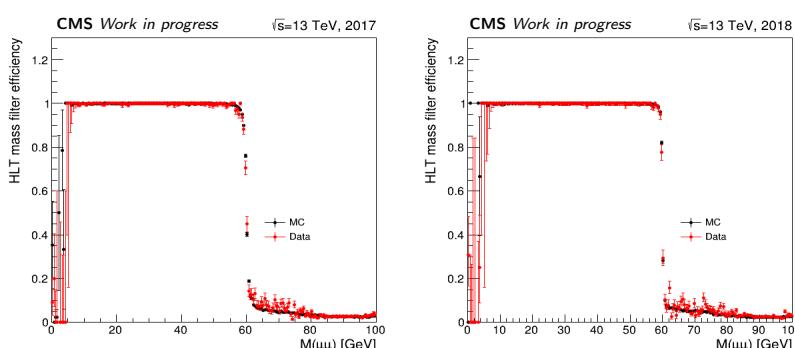


Figure 5.8: The dimuon invariant mass efficiency of the $\mu\mu + p_T^{\text{miss}}$ based trigger for 2017 (left) and 2018 (right).

Search for Supersymmetry with Soft Lepton Final States

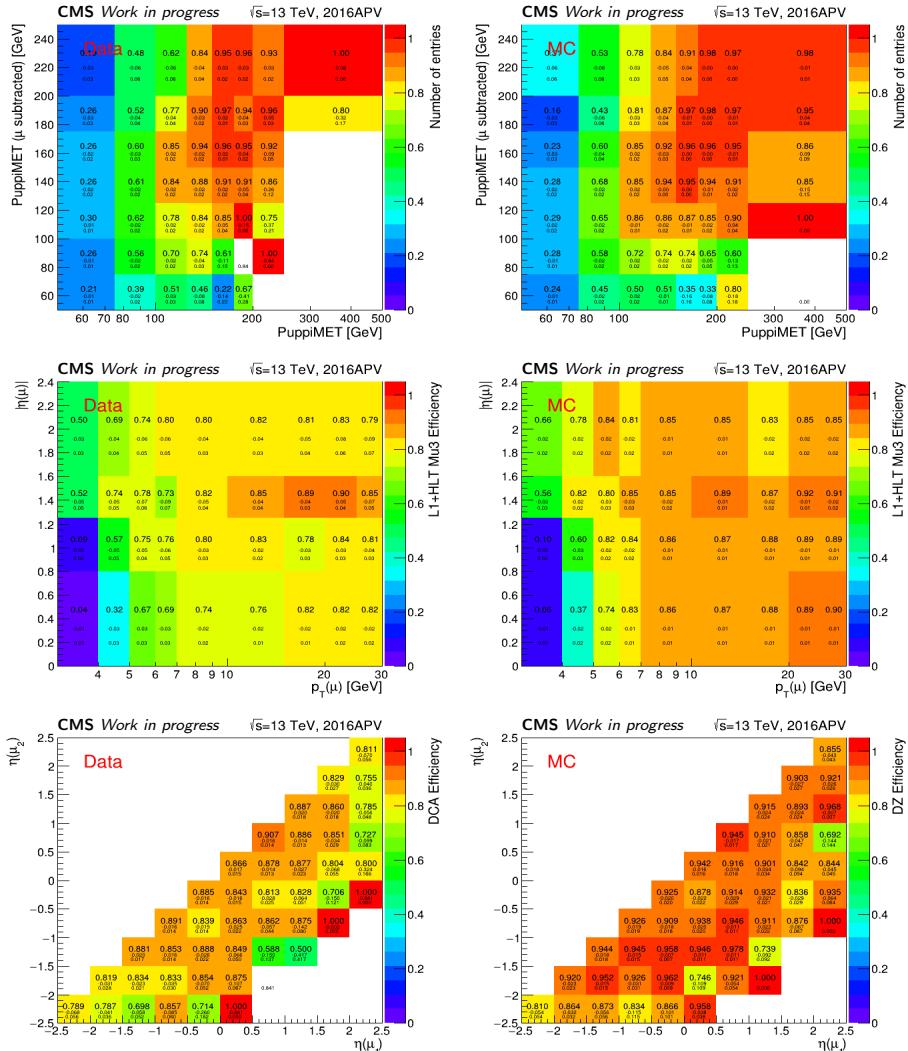


Figure 5.9: Efficiencies of the various HLT selections for data (left) and MC (right), corresponding to the 2016APV data-taking period. From top to bottom, the efficiency maps correspond to the p_T^{miss} , muon and DCA components of the $\mu\mu + p_T^{\text{miss}}$ based trigger, respectively.

5.4 Data and Simulated Samples

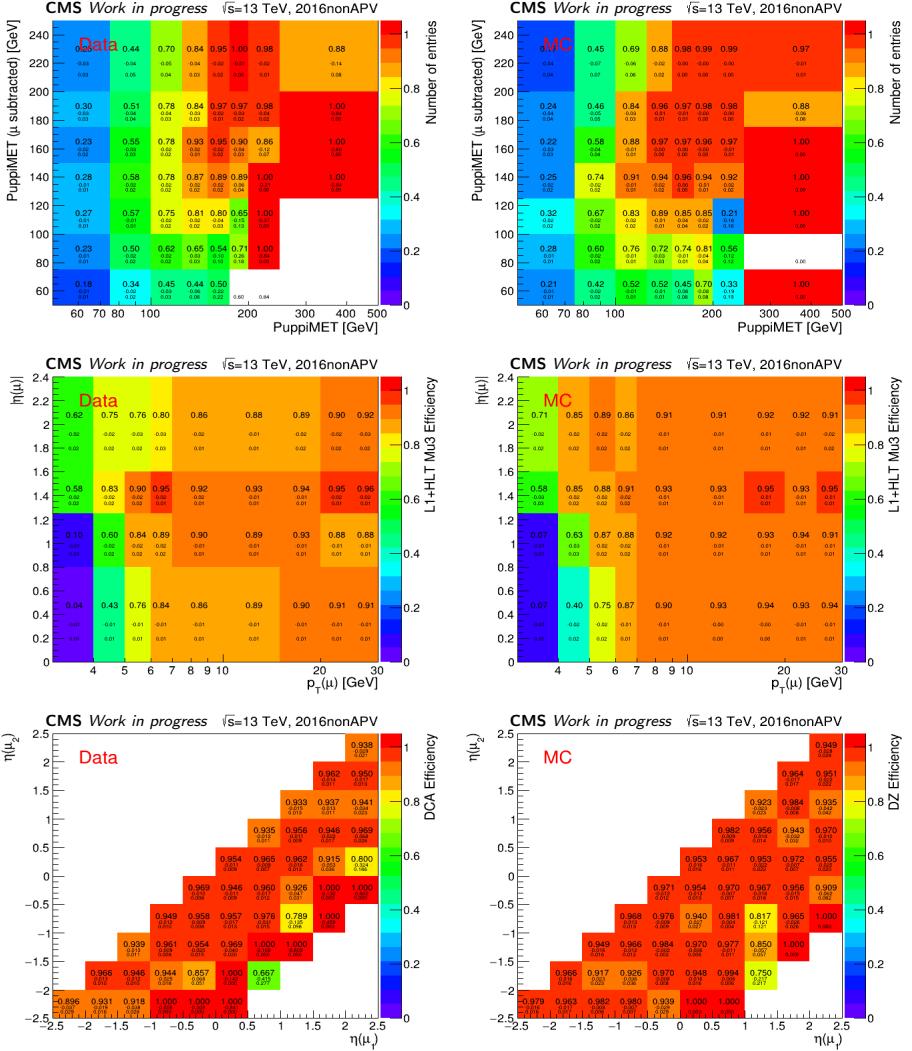


Figure 5.10: Efficiencies of the various HLT selections for data (left) and MC (right), corresponding to the 2016 data-taking period with HIP-mitigation. From top to bottom, the efficiency maps correspond to the p_T^{miss} , muon and DCA components of the $\mu\mu + p_T^{\text{miss}}$ based trigger, respectively.

Search for Supersymmetry with Soft Lepton Final States

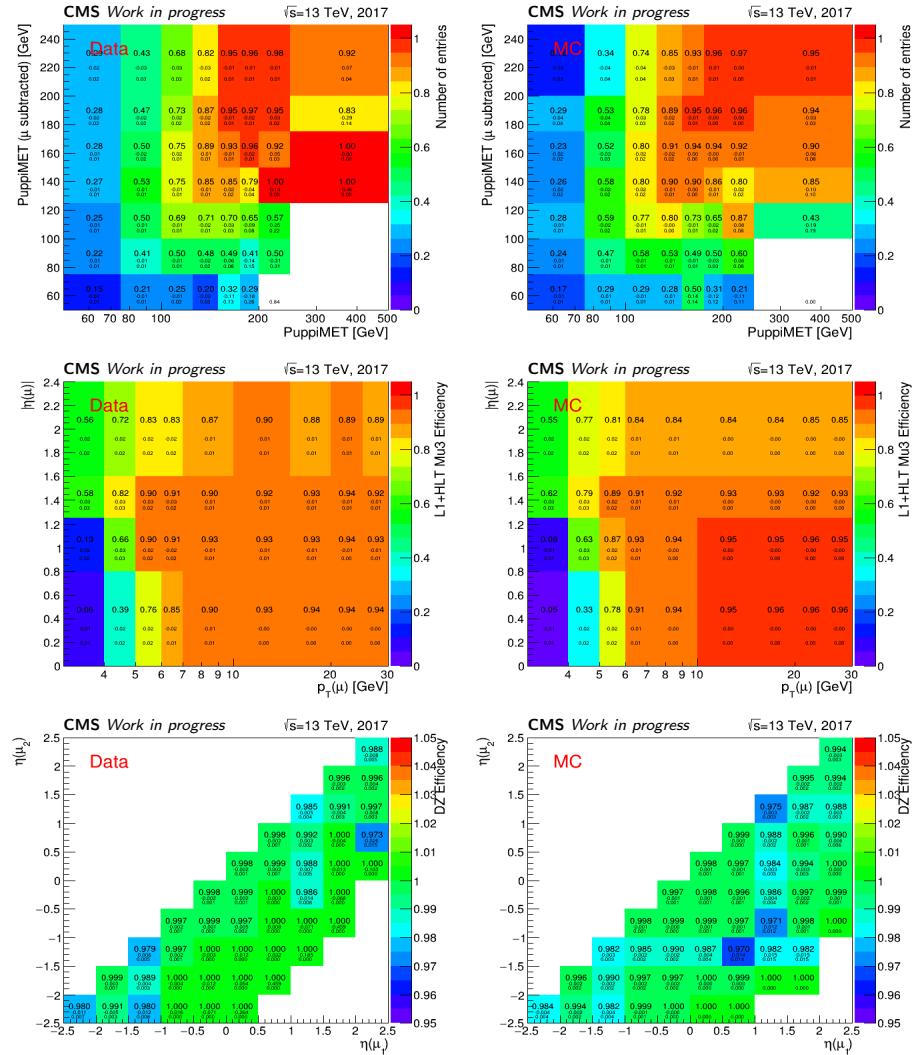


Figure 5.11: Efficiencies of the various HLT selections for data (left) and MC (right), corresponding to the 2017 data-taking period. From top to bottom, the efficiency maps correspond to the $P_{\text{T}}^{\text{miss}}$, muon and Δz components of the $\mu\mu + P_{\text{T}}^{\text{miss}}$ based trigger, respectively.

5.4 Data and Simulated Samples

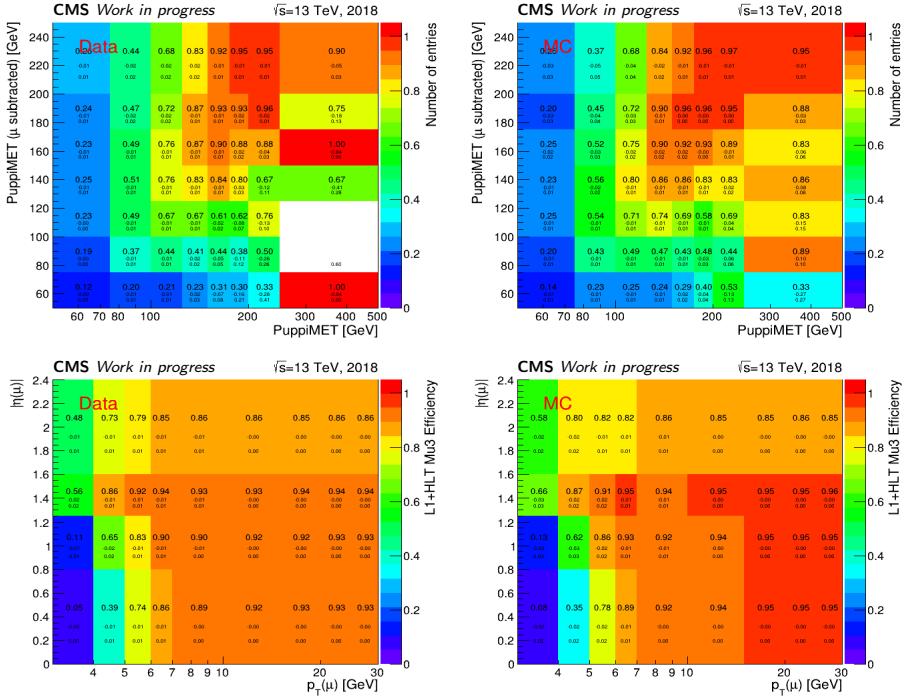


Figure 5.12: Efficiencies of the various HLT selections for data (left) and MC (right), corresponding to the 2018 data-taking period. The top (bottom) row corresponds to the p_T^{miss} (muon) component of the $\mu\mu + p_T^{\text{miss}}$ based trigger.

mass-splittings: $m_{\tilde{\chi}_2^0} (= m_{\tilde{\chi}_1^\pm})$ between 100 and 600 GeV in steps of 25 GeV, and $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ in (0.6, 0.8, 1, 1.5, 2, 3, 5, 7.5, 10, 15, 20, 30, 40, 50) GeV. For long-lived scenarios, samples are generated with the same mass-configuration as the prompt signal samples, but with $c\tau$ of $\tilde{\chi}_2^0$ in (0.1, 1, 10, 100) mm.

At the time of writing, signal samples for the higgsino scenario, corresponding to both topologies of Figure 5.1 and with $m_{\tilde{\chi}_1^\pm} = (\frac{1}{2}(m_{\tilde{\chi}_1^0} + m_{\tilde{\chi}_2^0})$, are still being produced. The remainder of this chapter will therefore focus on the wino-bino interpretation instead. The higgsino interpretation will be added in future.

Search for Supersymmetry with Soft Lepton Final States

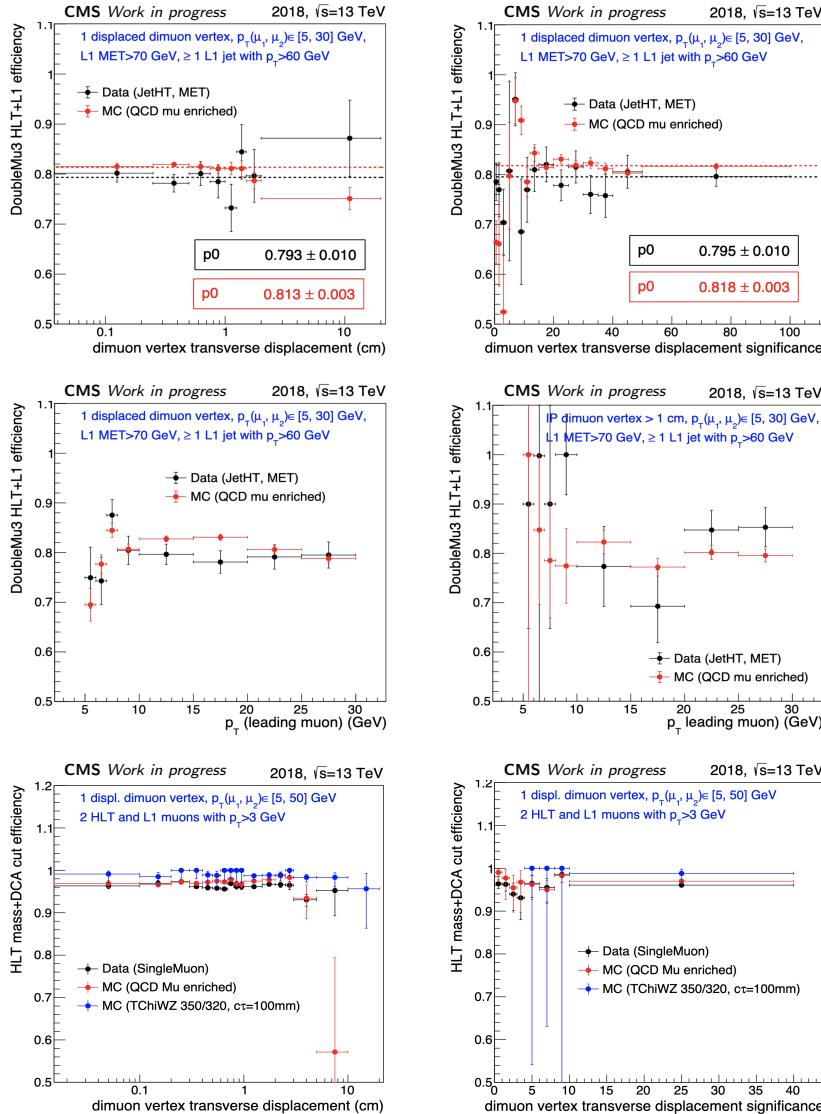


Figure 5.13: Efficiencies of the various muon-related HLT selections of the $\mu\mu + p_T^{\text{miss}}$ based trigger, corresponding to the 2018 data-taking period, as function of the transverse displacement of the dimuon vertex and the leading muon p_T . The first and second row correspond to the muon component, while the combined mass and DCA efficiency is shown in the third row.

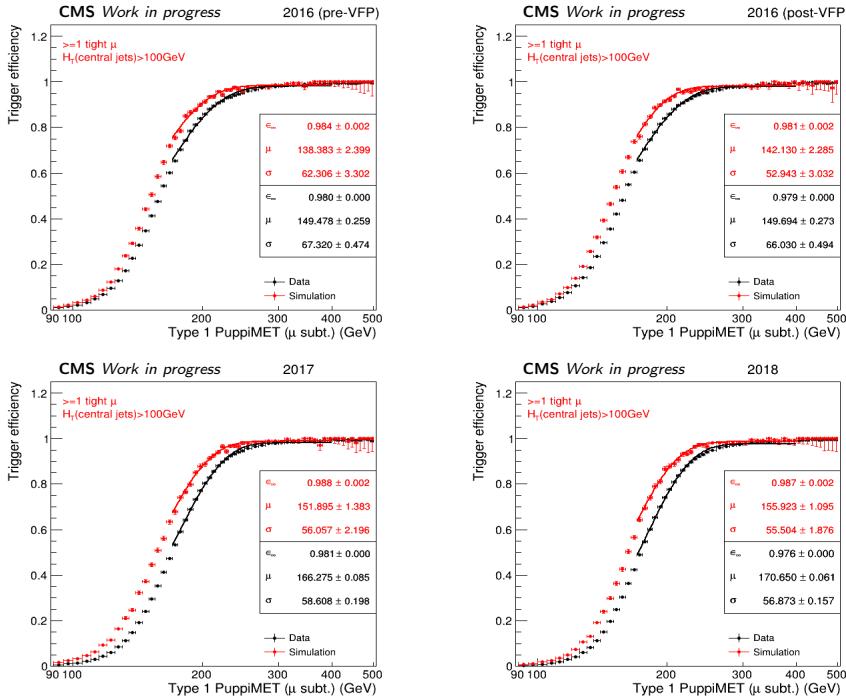


Figure 5.14: Efficiencies of the pure $p_T^{\text{miss,corr}}$ based trigger as function of offline p_T^{miss} for each data-taking period, measured in both data and MC simulation.

5.5 Object Definitions

Proton-proton (pp) collision events are reconstructed using the Particle Flow (PF) algorithm and the relevant physics objects are largely based on standard CMS reconstruction techniques, described in detail in Chapter 4. However, to target more compressed as well as long-lived signal scenarios, dedicated algorithms are utilized for the reconstruction of low- p_T electrons and displaced muon pairs. An overview of the physics object reconstructions is given below and more details are described in the respective sections.

- **Primary vertex** - The pp interaction vertex with highest quadratic sum of the p_T of associated jets and remaining tracks (Section 4.1.1), located within 24 cm in the z-direction and 2 cm in the transverse direction of the nominal interaction point.

- **Prompt electrons** (Section 5.5.1) - Two independent reconstructions are

used, being the standard GED electron (Section 4.3.4) and a new LowPt electron reconstruction (Section 4.3.5). The usage of LowPt electrons allows lowering the p_T threshold down to 1 GeV, which greatly increases the signal acceptance to more compressed signal scenarios (see Figure 5.7) compared to Ref. [89], which relied on GED electrons only with $p_T > 5$ GeV.

- **Prompt muons** (Section 5.5.2) - Both tracker muons and global muons (Section 4.2) are considered down to $p_T > 3.5$ GeV.
- **Displaced dimuons** (Section 5.5.3) - Pairs of standard muon tracks ($p_T > 3$ GeV) are refitted to a common displaced vertex without using the primary vertex as fit-constraint. This is a custom reconstruction procedure employed to target long-lived signal scenarios with two soft displaced muons from a common vertex, and was not considered in Ref. [89].
- **Jets and p_T^{miss}** (Section 5.5.1) - Standard charged-hadron-subtracted (CHS) PF jets (Section 5.5.4), reconstructed with the anti- k_T algorithm using distance parameter $R=0.4$. Jets from B-hadrons are tagged using the DeepJet algorithm. The transverse missing energy consists of the type-1 corrected p_T^{miss} based on PUPPI candidates (Section 4.6).

5.5.1 Prompt Electrons

The analysis relies on two independent electron reconstruction algorithms, being GED and LowPt electrons. The GED electron is the standard electron object used in the majority of CMS analyses and consists of a GSF track matched to an energy deposit in the ECAL. However, despite its good performance at high p_T , it is shown that the GED electron reconstruction efficiency drops rapidly below $p_T = 10$ GeV and is negligible at $p_T = 2$ GeV [175, 176]. The LowPt electron reconstruction consists also of matching a GSF track to an ECAL cluster, but instead relies on GSF tracks seeded by KF tracks down to $p_T = 1$ GeV. As such, it provides much higher efficiency for soft electrons and allows a great increase of the acceptance to highly compressed signal scenarios. The reconstruction of both types is described in more detail in Section 4.3.

In order to maximize the performance over the full p_T spectrum both the GED and LowPt electron collections are used. It should be noted however, that the reconstruction procedures are independent, hence it may happen that a given true electron is reconstructed as both types. To prevent double counting, a threshold

is defined at $p_T = 3$ GeV, marking the regimes for usage of LowPt and GED electrons: LowPt electrons are used for $1 < p_T < 3$ GeV and GED electrons are used for $p_T > 3$ GeV. This threshold - in combination with the custom identification working point described below - was found to yield the optimal trade-off between signal efficiency and background rejection over the full spectrum. The rare case of reconstructing a true electron both as a LowPt object with $p_T < 3$ GeV *and* as a GED object with $p_T > 3$ GeV is identified by finding pairs of objects satisfying $\Delta R(\text{LowPt}, \text{GED}) < 0.05$ and resolved by selecting the GED electron (discarding the LowPt electron) to avoid double counting. The reconstruction efficiency for the combined GED and LowPt electron collection is shown with the green curves in Figure 5.15.

After reconstruction, a series of selections are applied to define the electron objects used throughout the analysis. The selections are summarized in Table 5.2. The identification (ID) of electrons with $p_T > 10$ GeV relies on a standard multivariate (MVA) discriminator (`mvaFall17v2noIso`). To identify tight (loose) electrons³ with $p_T > 10$ GeV, the standard `WP90` (`WPL`) thresholds are used, which yield approximately 90 (98)% plateau efficiency with respect to reconstructed signal electrons. For electrons with $p_T < 10$ GeV, the identification strategy depends on the reconstruction type of the electron; For GED electrons (used down to $p_T = 3$ GeV), the identification relies on the same standard MVA, while for

Variable	Loose			Tight		
p_T	[1, 5)	[5, 10]	> 10	[1, 5)	[5, 10]	> 10
$ \eta $		≤ 2.5			≤ 2.5	
d_{xy} [cm]		< 0.05			< 0.05	
d_z [cm]		< 0.1			< 0.1	
IP_{3D} [cm]	< 0.05	$< 0.05 - 0.0065(p_T - 5)$	< 0.0175	< 0.025	$< 0.025 - 0.003(p_T - 5)$	< 0.01
$\sigma_{\text{IP}_{\text{3D}}}$	< 5	$< 5 - 0.5(p_T - 5)$	< 2.5	< 2.5	$< 2.5 - 0.1(p_T - 5)$	< 2
$PFIso_{\text{corr}, 0.3}^{\text{abs}}$ [GeV]		< 10			< 5	
DeepJet veto		–			< 0.45	
Electron MVA ID	custom "NewLoose" WP		WPL	custom "NewTight" WP		WP90
No missing pixel hits		✓			✓	
Conversion veto		✓			✓	

Table 5.2: List of selection criteria imposed on reconstructed electron candidates to define loose and tight electrons. The custom "NewTight" and "NewLoose" working points correspond to the standard `mvaFall17V2noIso` discriminator for GED Electrons, and the dedicated (2020Nov28) discriminator for LowPt Electrons. They are defined such as to maintain at low p_T the same efficiency as the centrally derived `WP90` (`WPL`) working points at the plateau with respect to reconstructed signal electrons, being approximately 90 (98)%

³Tight electrons are the main electron objects used throughout the analysis, while loose electrons are defined with relaxed selection and used for the fake lepton background estimate with the tight-to-loose method described in Section 5.7.

LowPt electrons (used for $1 < p_T < 3$ GeV) the identification relies on a dedicated LowPt electron discriminator (2020Nov28). Custom thresholds (also referred to as “working points”) are designed such that below $p_T = 5$ GeV the identification efficiency is similar to that of the standard WP90 and WPL thresholds at the plateau, while for $5 < p_T < 10$ GeV a smooth efficiency transition is obtained. The ID efficiencies are shown with the orange curves in Figure 5.15.

Isolation criteria are imposed to reduce (fake lepton) backgrounds from hadronic activity. These criteria are based on the standard PF isolation, defined as:

$$PFIso_{0.3}^{abs} = \sum_{h^0} p_T^{h^0} + \sum_{\gamma} p_T^{\gamma} + \sum_{h^{\pm}} p_T^{h^{\pm}} \quad (5.5)$$

where h^0 , γ and h^{\pm} denote PF neutral hadrons, PF photons and PF charged hadrons, respectively, in the vicinity of the reference electron (in this case within $\Delta R < 0.3$). Leptons are not included in the computation. A subtlety in this computation arises due to the fact that leptons in the analysis regions (shown in Section 5.6) are allowed to be close-by, down to $\Delta R = 0.05$, and therefore enter each others isolation cone, combined with the fact that electrons are not necessarily PF electrons. The latter is true since the *LowPt Electron* reconstruction is independent from the GED (and thereby PF) reconstruction, while GED electrons themselves in this analysis are not explicitly required to pass the PF ID. It may therefore happen that a given electron (either a *LowPt Electron* or a GED electron failing PF ID) is also reconstructed as a PF charged hadron, and therefore is included in the isolation sum of another nearby electron. As a result, such cases would make two closeby electrons (eg. from signal) appear less isolated, potentially failing the isolation criteria. To prevent such scenarios and have a consistent isolation computation (that does not include leptons), the above isolation computation is corrected in order to exclude non-PF electrons from the isolation sum of other close-by electrons. The procedure is as follows; Assume a reference electron (of any type) that has another electron within its isolation cone. A correction is needed depending on which of the four following situations applies:

- The second electron is a GED electron that passes PF ID. No correction is necessary, since it is already excluded from the isolation computation of Equation 5.5.
- The second electron is a GED electron that fails PF ID. Its components are then assumed to be used in the reconstruction of another PF candidate (eg. PF charged hadron) and a correction is needed.

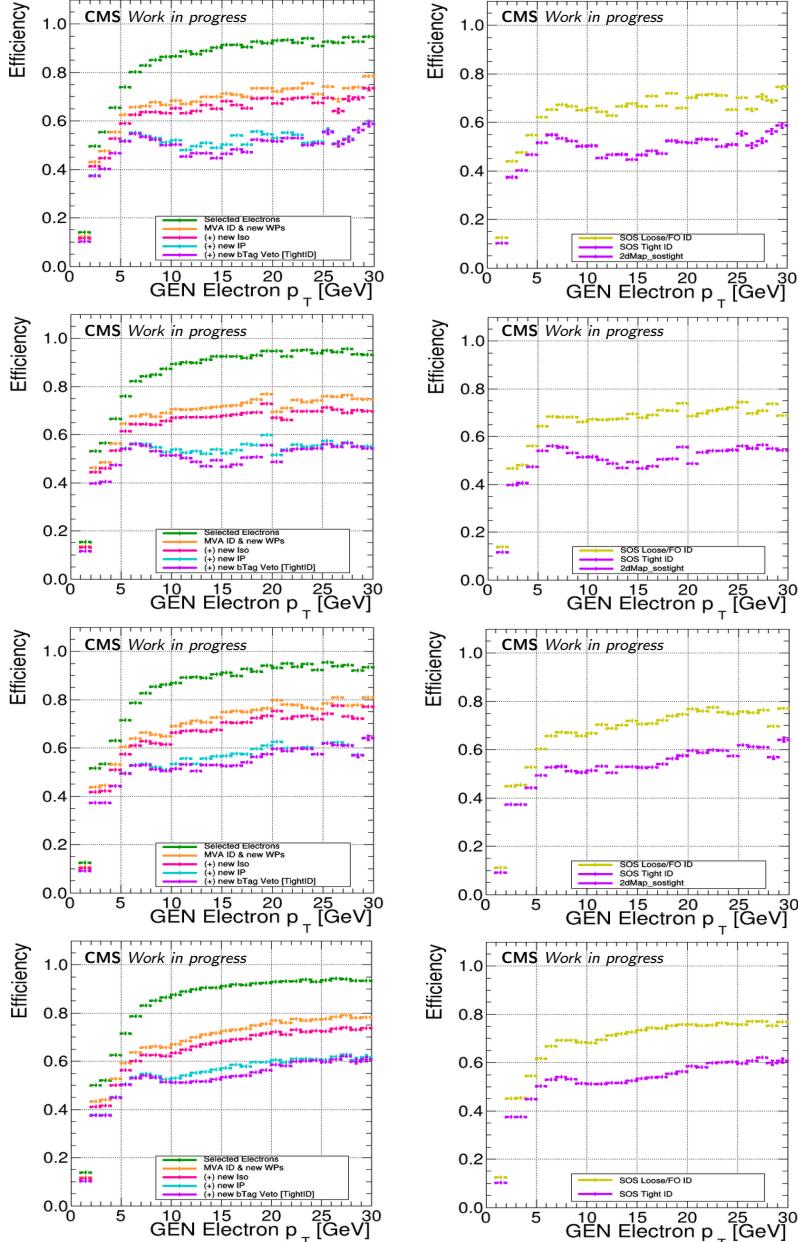


Figure 5.15: Electron efficiencies corresponding to the selections of Table 5.2. From top to bottom: 2016APV, 2016, 2017, 2018. Left: Efficiencies factorized in five steps. The green curve corresponds to the reconstructed electrons (LowPt for $p_T < 3$ GeV and GED for $p_T > 3$ GeV) after overlap removal, while the purple curve corresponds to the tight electron definition of Table 5.2. The intermediate steps are the applied ID (orange), isolation (pink), and the 3D impact parameter and its significance (cyan). Right: Efficiency for the tight (purple, same as in the left column) and loose (yellow) electron definitions.

- The second electron is a *LowPt Electron* that overlaps (within $\Delta R < 0.05$) with a GED electron that passes PF ID. It is then assumed to be the same particle, and since the GED electron already passes PF (same as the first bullet), no correction is needed.
- The second electron is a *LowPt Electron* that does not overlap with a GED electron passing PF ID. Its components are assumed to be used in the reconstruction of another PF candidate (eg. PF charged hadron) and a correction is needed.

If a correction is needed, then the p_T of the KF track associated to the second electron is subtracted from the PF isolation of the reference electron. The p_T of the KF track drives the momentum assigned to charged PF hadrons and thus is used (as opposed to the associated GSF track-driven p_T) in order to consistently cancel out the isolation term from PF candidates that are not electrons. Furthermore, for a fully consistent treatment, the above correction would in principle also have to be applied to muons that have a closeby non-PF electron. However, because this topology in signal is extremely rare (it requires a non-PF electron and a muon - one from the neutralino decay, the other from the chargino decay - that accidentally are geometrically closeby) this is not applied. Lastly, electrons by default can never contribute to their own isolation sum; For GED electrons, the PF algorithm keeps track of which elements (tracks/clusters) are shared between GED electrons and other PF candidates via a particle mapping and as such excludes these candidates from the isolation of the GED electron. *LowPt Electrons* are excluded from their own isolation sum via a standard *dead-cone* definition that disallows contributions from very-nearby particles. The efficiency of the isolation criteria (Table 5.2) is shown with the pink curves in Figure 5.15.

Finally, requirements on the 3-dimensional impact parameter (significance) IP_{3D} ($\sigma_{\text{IP}_{3D}}$) are applied to reject non-prompt lepton backgrounds. Electrons compatible with B-hadron decays are additionally vetoed by evaluating the DeepJet b tagger for jets associated to electrons (by common ancestor for GED and by ΔR for LowPt electrons). An upper bound on the tagger score is chosen by requiring minimal flavor dependency of the fake rate (see Section 5.7), while maintaining efficiency. The associated efficiencies are shown with the blue and purple curves in Figure 5.15.

Scale factors

The efficiency of prompt electrons to pass the tight selection predicted by MC simulation is corrected in order to match that observed in data. To facilitate these measurements, the total efficiency can be broken into several sub-components as:

$$\epsilon(e) = \epsilon(e)_{\text{Reco}} * \epsilon(e)_{\text{Tight ID+Iso} \mid \text{Reco}} * \epsilon(e)_{\text{Tight IP cuts} \mid \text{Reco} + \text{Tight ID+Iso}} \quad (5.6)$$

where the first term is the reconstruction efficiency and the second term corresponds to the efficiency of a reconstructed object to also pass the tight identification requirements (MVA discriminant cuts, conversion veto, and missing pixel hits) and isolation described above. The last term is the efficiency of a prompt electron to additionally satisfy requirements on the quality of the associated track-to-vertex match (impact parameters and B veto). Efficiencies are measured in both data and MC using the tag-and-probe method, utilizing the two-prong decay of a known resonance. The MC efficiency is calculated by fitting a functional form to events with prompt leptons that are truth-matched to the target resonance. The data efficiency is calculated by fitting a signal (parameterized identically to the MC case above) plus a smooth background to the data distribution.

Reconstruction scale-factors (SF) are provided centrally by CMS for GED electrons with p_T as low as 10 GeV. Those measurements utilize Z boson decays and are performed separately in bins of probe electron p_T and η . The SFs are very close to unity for all years and detector regions, except for the region $1.44 < |\eta| < 1.57$, which corresponds to the crack between the ECAL barrel and end-caps. For $p_T < 10$ GeV, dedicated SFs are measured with respect to tracks using $B \rightarrow K^\pm J/\psi(ee)$ candidates. A third track leading to a mass consistent with B decays is required to increase the candidate purity, and all tracks must be separated by $\Delta R > 0.3$. Tag electrons are the tracks qualifying as GED electron with $p_T > 3$ GeV and passing the standard MVA ID with an 80% working point, while Kaon candidates must have $p_T > 5$ GeV and $\Delta R > 0.1$ from the probe track. The SFs are extracted in bins of probe p_T via fits to the J/ψ mass, an example of which is shown in Figure 5.16. The LowPt electron SFs amount to approximately 1 ± 0.1 for $1 < p_T < 5$ GeV and 0.98 ± 0.05 for $5 < p_T < 10$ GeV. For GED electrons, the SFs are 0.91 ± 0.12 for $2 < p_T < 5$ GeV and 0.92 ± 0.08 for $5 < p_T < 10$ GeV.

The remaining components of the tight electron efficiency are measured using a combination of Z and J/ψ decays. While the Z provides an excellent source of prompt leptons at moderate p_T , the measurement becomes more challenging

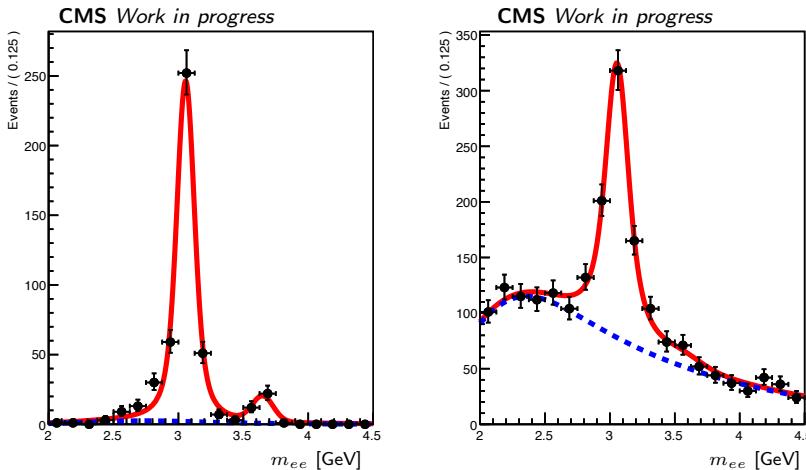


Figure 5.16: Example fits to the J/ψ and $\psi(2S)$ mass distributions for passing (left) and failing (right) probes in the tag-and-probe measurement of the reconstruction efficiency for GED electrons with p_T from 2 to 5 GeV.

below 5-10 GeV due to decreasing purity. For this phase-space, the J/ψ becomes attractive, since it is copiously produced at the LHC and typically decays to softer leptons. The tight electron scale factors for the full p_T spectrum are shown in Figure 5.17.

For $p_T > 10$ GeV, the complete tight selection efficiency with respect to all reconstructed objects is measured solely with standard Drell-Yan events, separately for the barrel and end-cap regions. Events are selected using an un-prescaled single-electron trigger (HLT_Ele32_WPTight_Gsf), with tag electrons required to have $p_T > 34$ GeV, $|\eta| < 2.5$, pass a cut-based ID requirement and match the corresponding trigger object. Probes must have opposite charge, be well-separated from the tag ($\Delta R > 0.3$) and give rise to an invariant mass $m_{\ell\ell}$ from 60 to 140 GeV.

For $p_T < 10$ GeV, the ID and isolation SFs are measured with J/ψ candidates. It should be noted, that J/ψ are mostly produced through pure QCD interactions and yield low- p_T leptons, which would not pass the thresholds of single-electron triggers. Therefore, J/ψ candidates from *pileup* interactions are collected parasitically in events that independently fired the single lepton trigger due to the presence of eg. a lepton from W boson production in the primary pp interaction. As such, neither tag nor probe are matched to the triggering object. The tag electron is required to have $p_T > 3$ GeV and satisfy the tight selection criteria of Table 5.2. The probe electron should have opposite sign form a dilepton mass

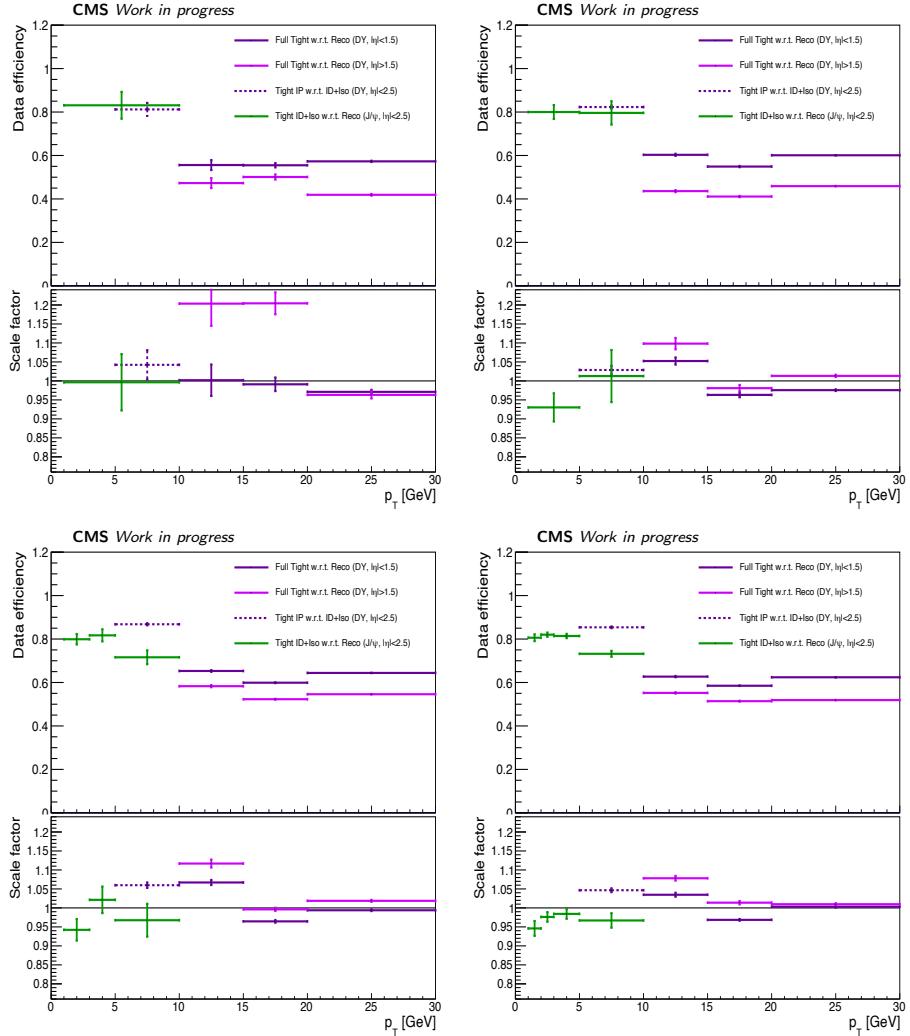


Figure 5.17: Measurements of the various components of the total electron scale factor, as obtained with samples of DY and J/ψ decays in bins of probe p_T and η . In the legend, “ID” refers to the requirements on the MVA ID, pixel hits and conversion veto. “Iso” refers to the isolation requirement and “IP” refers to the impact parameter cuts and DeepJet veto. “Full Tight” refers collectively to all selections used to define tight electrons (Table 5.2). Results are shown for 2016 data periods, before (top left) and after (top right) HIP-mitigation, 2017 (bottom left), and 2018 (bottom right).

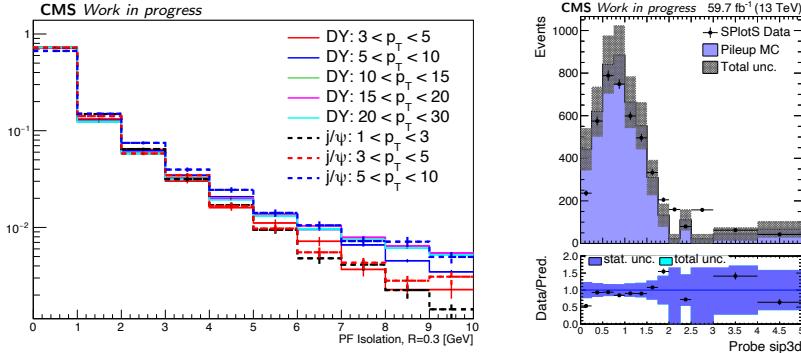


Figure 5.18: Left: Isolation sums for electrons entering the Drell-Yan and J/ψ SF measurement regions in bins of p_T in MC events, normalized to the total number of events in each sample. Right: Three dimensional impact parameter significance σ_{IP3D} for probe electrons in the J/ψ SF measurement region. The data points are based on fits to the $m_{\ell\ell}$ spectrum of the tag and probe pair in each bin of σ_{IP3D} , to unfold the contribution from J/ψ candidates.

from 2.4 to 3.8 GeV. The relative distance in z between the tag and probe leptons should be less than 1 mm, but greater than 1 mm from the PV, and should match to the same (non-PV) vertex. It may furthermore be noted that J/ψ candidates from pileup can be used not only for the measurement of the ID SF, but also for the isolation SFs, since they are typically produced at rest, hence the decay products are approximately back-to-back and well-isolated. This is shown by the excellent agreement of the isolation sums for electrons from Drell-Yan and J/ψ in Figure 5.18 (left).

Finally, for $5 < p_T < 10$ GeV, the SFs related to the impact parameters are measured with Z decays using a sample of leptons that already satisfy the tight ID and isolation, given the large background for failing probes when measuring with respect to all reconstructed candidates instead. For $1 < p_T < 5$ GeV, there are not sufficient candidates from DY, nor can the J/ψ resonance be used, since some population of J/ψ are produced in B decays and are thus not necessarily prompt. However, Figure 5.18 (right) proves that the impact parameter is well modeled by MC simulation, hence a SF of unity is applied for these electrons instead, with a conservative uncertainty of 20%.

5.5.2 Prompt Muons

The analysis relies on standard muon reconstruction techniques as described in Section 4.2. Tracker and global muons are used down to $p_T = 3.5$ GeV, both

of which consist of a Tracker track matched to at least one track segment in the muon chambers. Standalone muons, reconstructed with information from the muon subdetectors only, are not used due to their relatively poor momentum resolution and decreased efficiency for low p_T objects.

The selection criteria for reconstructed muons are shown in Table 5.3. The identification requirements are based on centrally provided selections and are largely the same as in Ref [89]. In particular, muons must pass the *loose ID*, which demands that tracker or global muons are additionally identified as PF muon, as well as the *soft ID*, which is optimized for low p_T muons and ensures a high purity via tight track segment matching requirements [130]. Tight cuts on isolation, impact parameters and B-hadron compatibility are applied to reject muons from (non-prompt) hadron decays. To veto muons compatible with B decays, the DeepJet tagger is used instead of the DeepCSV tagger (used in Ref. [89]) due to its generally higher performance. The threshold on the tagger score is chosen to minimize the flavor dependency of the muon fake rate, while maintaining high signal efficiency. The muon efficiencies are shown in Figure 5.19.

Variable	Loose	Tight
p_T [GeV]	≥ 3.5	≥ 3.5
$ \eta $	≤ 2.4	≤ 2.4
loose ID	✓	✓
soft ID	✓	✓
$PFIso_{0.3}^{abs}$ [GeV]	–	< 5
$PFIso_{0.3}^{rel}$	< 1.0	< 0.5
IP_{3D} [cm]	< 0.0175	< 0.01
$\sigma_{IP_{3D}}$	< 2.5	< 2.0
d_{xy} [cm]	< 0.05	< 0.05
d_z [cm]	< 0.1	< 0.1
DeepJet veto	–	< 0.15

Table 5.3: List of selection criteria imposed on reconstructed muon candidates to define loose and tight muons.

Scale factors

To correct the efficiency of the prompt muon reconstruction and selection in MC simulated events to that in observed data, SFs are derived and applied to the simulated events. To facilitate the SF measurements, the muon efficiency is factorized

Search for Supersymmetry with Soft Lepton Final States

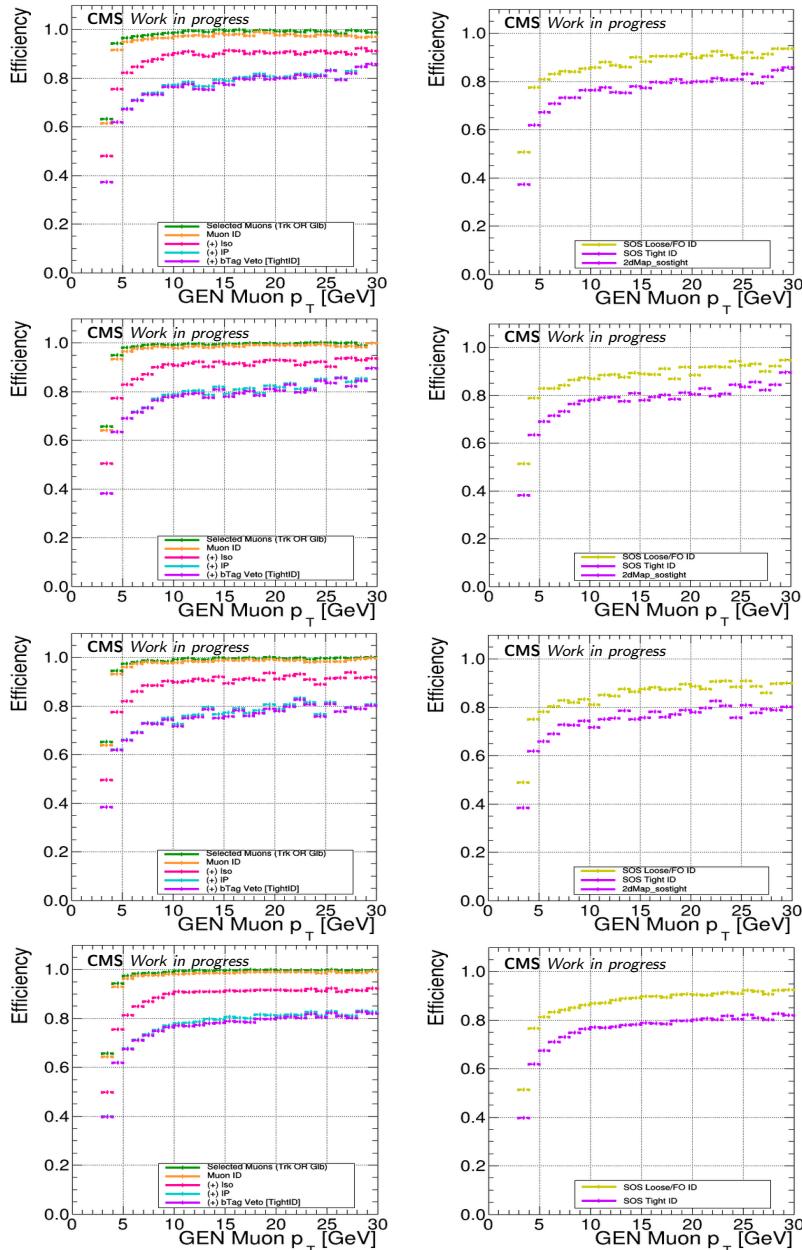


Figure 5.19: Muon efficiencies corresponding to the selections of Table 5.3. From top to bottom: 2016APV, 2016, 2017, 2018. Left: Efficiencies factorized in five steps. The green curve corresponds to the reconstructed muons, while the purple curve corresponds to the tight muon definition of Table 5.3. The intermediate steps are the applied IDs (orange), isolation (pink), and the 3D impact parameter and its significance (cyan). Right: Efficiency for the tight (purple, same as in the left column) and loose (yellow) muon definitions.

into four terms:

$$\begin{aligned}\epsilon(\mu) = & \epsilon(\mu)_{\text{TrackRECO}} * \epsilon(\mu)_{\text{Muon} \mid \text{TrackRECO}} * \epsilon(\mu)_{\text{POG Loose ID} \mid \text{Muon}} \\ & * \epsilon(\mu)_{\text{Tight ID} \mid \text{POG Loose ID}}\end{aligned}\quad (5.7)$$

The first term ($\epsilon(\mu)_{\text{TrackRECO}}$) corresponds to the track reconstruction and was found to be very close to unity in both data and MC, hence can safely be neglected. The second term ($\epsilon(\mu)_{\text{Muon} \mid \text{TrackRECO}}$ and third term ($\epsilon(\mu)_{\text{POG Loose ID} \mid \text{Muon}}$) account for the efficiency of reconstructed tracks to pass the standard muon *loose ID*. Both terms are provided centrally by CMS and are based with tag-and-probe measurements on Z or J/ψ resonances in bins of probe muon p_T and η . The SFs are very close to unity as well, except for very soft muons ($p_T \simeq 3$ GeV) in the barrel region, since those hardly reach the muon chambers due to the strong track curvature and are more difficult to reconstruct.

The last term ($\epsilon(\mu)_{\text{Tight ID} \mid \text{POG Loose ID}}$) denotes the efficiency for muons that pass the loose ID to additionally pass the remaining tight identification criteria of Table 5.3. It is measured with the same tag-and-probe technique on a sample of muons from Z decays, selected with an un-prescaled single muon trigger (`HLT_IsoMu24`). The tag muon is required to be matched to the triggering muon, pass $p_T > 29$ GeV, $|\eta| < 2.4$ and a centrally provided tight identification criterion [130]. The probe muon must have opposite charge, $p_T > 3$ GeV, $\eta < 2.4$, pass the loose ID and yield a dimuon mass loosely compatible with a Z boson ($60 < m_{\ell\ell} < 120$ GeV). In bins of probe p_T and η , the $m_{\ell\ell}$ distributions are fit with a functional form (separately for probes passing and failing the tight selections of Table 5.3), and efficiencies are calculated by dividing the resulting amount of passing probe muons by the total amount of probe muons from Z . The efficiencies and SFs are shown in Figure 5.20.

5.5.3 Displaced Dimuons

To target signal scenarios where the $\tilde{\chi}_2^0$ is long-lived, the analysis includes a dedicated search region category that is based on the presence of a high quality soft displaced muon pair. The reconstruction of these objects is based on a custom procedure that consists of fitting pairs of standard muon tracks to a common (displaced) vertex without using the primary pp interaction vertex as fit constraint. The track fitting relies on the usual Kalman filter formalism [151–153] and is mathematically equivalent to a global least-squares minimization. Given the targeted displacement regime of the analysis (up to roughly 10 cm) and the usage of tracker

Search for Supersymmetry with Soft Lepton Final States

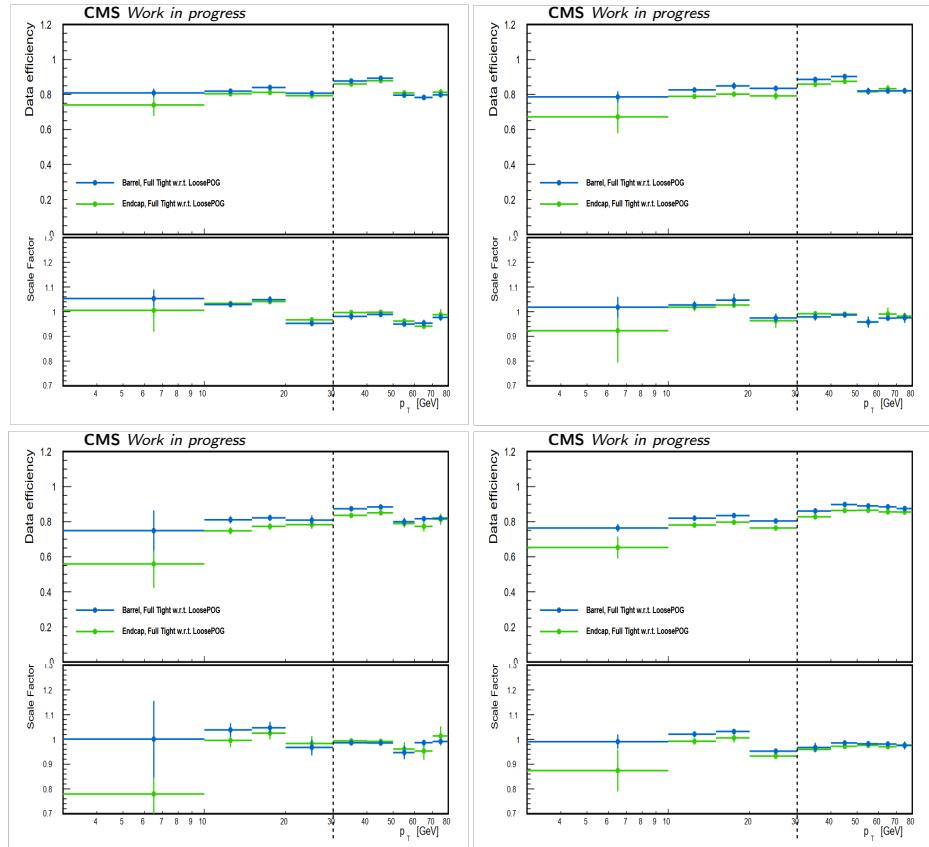


Figure 5.20: Measurements of the muon scale-factors corresponding to the last term of Equation 5.7 ($\epsilon(\mu)_{\text{Tight ID} \mid \text{POG Loose ID}}$), as obtained in samples of Z boson decays, in bins of probe p_T and η . Results are shown for 2016 data periods, before (top left) and after (top right) HIP-mitigation, 2017 (bottom left), and 2018 (bottom right). The dashed line at 30 GeV indicates the upper bound of the p_T of muons used in the analysis.

and global muons, the position of the newly constructed secondary vertex must be located within the Tracker volume. The parameters of the muon tracks are constrained to the displaced vertex in order to improve the resolution of kinematic variables. The resolution of the 2D vertex position in the transverse plane was found to be approximately 2.5 ± 0.1 mm for displacements up to 10 cm, which increases to 4.5 ± 0.3 mm above 10 cm due to the decreasing number tracker layer hits towards higher vertex displacements.

Table 5.4 shows the list of selection criteria applied on reconstructed dimuon candidates in order to be identified as a tight dimuon. Selections are imposed both on the individual dimuon legs, and on the refitted vertex itself; The identification criteria imposed on the dimuon legs are essentially the standard *loose ID* and *soft ID* that provided centrally and used also for prompt muons, except that requirements on the impact parameter and the number of pixel and tracker hits are removed, in order to maintain efficiency for candidates with larger displacement.

Variable	Loose	Tight
p_T [GeV]	> 3.0	> 3.0
$ \eta $	< 2.4	< 2.4
Charge	OS	OS
IP_{3D} [cm]	> 0.0175	> 0.0175
Displaced ID	PF and (Global or Tracker) —	PF and (Global or Tracker) highPurity and TMOneStationTight
$PFIso_{0.3}^{abs}$ [GeV]	< 50	< 10
$PFIso_{0.3}^{rel}$	—	< 1
B vertex veto	—	✓
L_{xy} [cm]	> 0.05	> 0.05
DCA _{3D} [cm]	< 0.2	< 0.2
Fit probability	> 0.1	> 0.1

Table 5.4: List of selection criteria imposed on reconstructed dimuon candidates to define loose and tight dimuons. The first block of selections applies the individual dimuon legs (ie. the muons that make up the pair), while the second block of selections applies to the vertex (fit).

The isolation criteria rely on the standard PF isolation variables; It should be noted that the charged component of the PF isolation sum (Equation 5.5) only accounts for charged hadrons that are associated to the PV. As such, displaced dimuons from background processes (eg. non-prompt B decays from the hard interaction or tracks from pileup) could *appear* more isolated in cases where nearby

charged hadrons are not matched to the PV⁴. However, it was found that the standard PF isolation still works well in selecting signal and rejecting background candidates, even at large displacements. The isolation thresholds themselves are slightly relaxed compared to those of tight prompt muons (Table 5.3) in order to maximize signal efficiency while accepting only a small increase in background events.

The 3D distance of closest approach (DCA) at any point along the muon tracks as well as the fit probability cuts imposed on the reconstructed vertex serve to retain only good quality muon pairs, and reject muons that are far away and incompatible with coming from the same vertex. Lastly, lower bounds on the impact parameters are applied both on the legs (IP_{3D}) and the vertex (L_{xy}) in order to have orthogonal object selection with respect to the prompt muons as defined in Table 5.3. The reconstruction efficiency and the efficiency of the additional identification and isolation criteria for displaced muon pairs are shown as function of displacement in Figure 5.21.

Scale factors

The efficiency of displaced dimuon objects to pass the tight selection (Table 5.4) predicted by MC simulation is corrected in order to match that observed in data. To facilitate these SF measurements, the total efficiency is factorized as:

$$\epsilon(\mu\mu) = \epsilon(\text{SV, tracks})_{\text{Reco}} * \epsilon(\mu\mu)_{\text{Tight Dimuon ID} \mid \text{Reco}} \quad (5.8)$$

The first term ($\epsilon(\text{SV, tracks})_{\text{Reco}}$) corresponds to the reconstruction efficiency of the two displaced tracks as well as the efficiency of reconstructing the secondary vertex given the two tracks. The associated correction factors are obtained by measuring (displaced) $K_s^0 \rightarrow \pi^\pm \pi^\mp$ decays, which resembles the signal topology of the present search - ie. a neutral particle decaying within the tracker volume to two oppositely charged tracks. To obtain an unbiased sample of K_s^0 candidates, $Z \rightarrow \mu\mu$ events are selected with a mixture of double-muon triggers, while the K_s^0 candidates are reconstructed in additional jets and hadronic activity in the underlying event. The triggering muons are required to satisfy a standardized muon

⁴This in principle also affects the neutral component, which includes all neutral hadrons and photons in the isolation cone. This contribution is typically corrected to account for candidates from pileup, by subtracting the energy deposits from charged particles within the isolation cone not associated to the PV, multiplied by a factor 0.5 (approximately the ratio of neutral to charged hadron production in the hadronization process of pileup interaction) [226]. At large displacements PF candidates tend to be not associated to the PV, hence the correction increases and the reference candidate would appear more isolated.

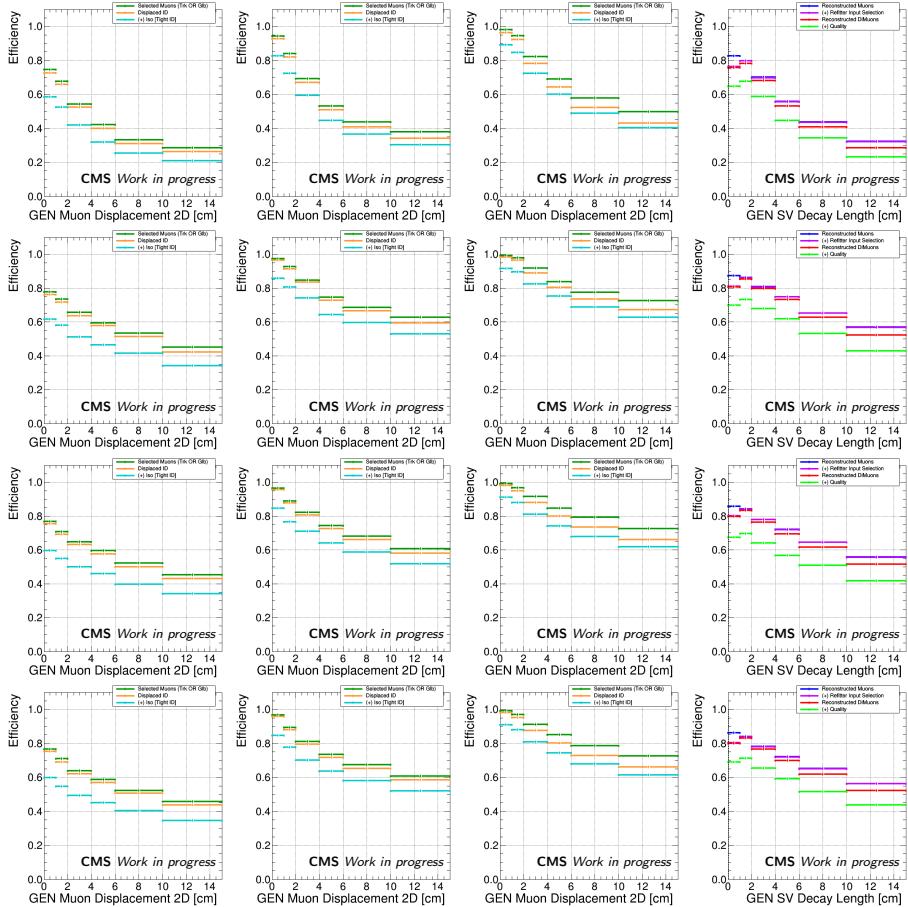


Figure 5.21: Efficiencies of reconstructed muons as function of displacement. The first three columns display the reconstruction, identification and isolation efficiency for individual muons satisfying $3 < p_T < 5$, $5 < p_T < 10$ and $10 < p_T < 30$ GeV. The right-most column shows the efficiency of reconstructed muon pairs as well as the quality (DCA and fit-probability) requirements. From top to bottom: 2016APV, 2016, 2017, 2018.

identification (*medium ID*), $p_{T,\mu_1}(p_{T,\mu_2}) > 30$ (25) GeV and yield a dimuon mass within 10 GeV of the Z resonance. Events with loosely tagged b jets are vetoed. Kaon candidates are reconstructed by refitting pairs of oppositely charged tracks with $p_T > 1$ GeV and $\text{DCA} < 0.2$ cm to a common high-quality ($\chi^2 < 7$) vertex of which the position in the transverse plane with respect to the beamline is collinear with the di-track momentum. The di-track mass must be within 70 MeV of the K_s^0 mass. The background contamination from coincidentally compatible tracks is subtracted with a linear fit to the candidates outside the kaon mass-window. The scale factors are then computed as the ratio of selected kaon candidates in observed data and MC simulated events, measured separately in bins of the p_T of the track pair and transverse displacement. To factor out overall K_s^0 mismodeling effects - assumed to be independent on the radial decay distance - and thereby isolate the systematic discrepancy between simulation and data as function of the radial position of the vertex, the number of kaon candidates in simulation is normalized to data using only the candidates in the first bin of displacement (up to 0.5 cm). The resulting scale factors to be applied per (dimuon) object are measured for the full Run-2 data taking period and are shown in Figure 5.22.

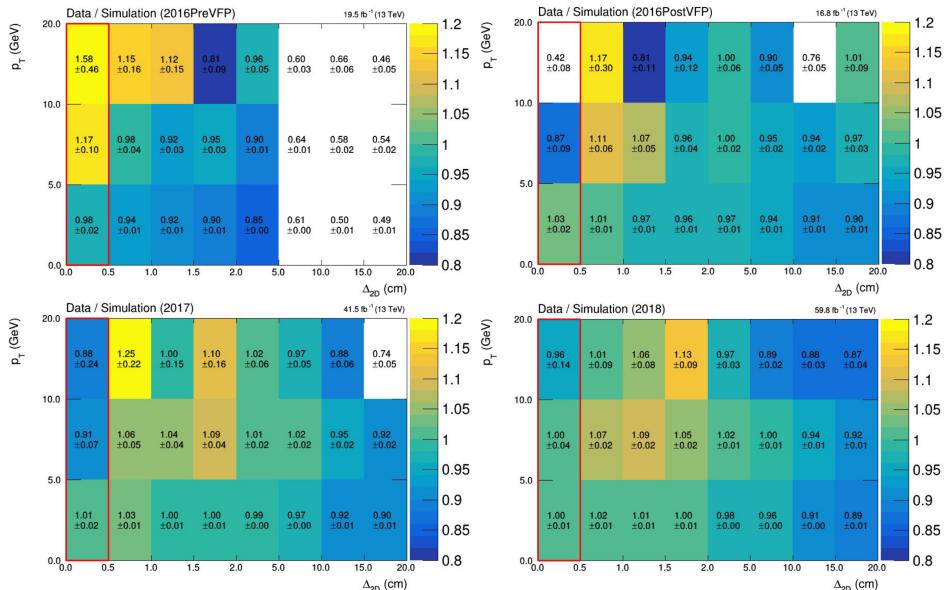


Figure 5.22: Correction factors for the di-track and vertex reconstruction efficiency, measured with $K_s^0 \rightarrow \pi\pi$ decays in $Z+jets$ events, and binned in dimuon p_T and radial distance to the PV. Results are shown for 2016 data periods, before (top left) and after (top right) HIP-mitigation, 2017 (bottom left), and 2018 (bottom right).

For the second term of Equation 5.8 ($\epsilon(\mu\mu)_{\text{Tight Dimuon ID} \mid \text{Reco}}$) a per-muon correction is applied to both legs of the dimuon object:

$$\epsilon(\mu\mu)_{\text{Tight Dimuon} \mid \text{Reco}} = \epsilon(\mu_1)_{\text{Tight Displaced ID} \mid \text{Reco}} * \epsilon(\mu_2)_{\text{Tight Displaced ID} \mid \text{Reco}} \quad (5.9)$$

The per-muon efficiency is further factorized into three terms:

$$\begin{aligned} \epsilon(\mu)_{\text{Tight Displaced ID} \mid \text{Reco}} &= \epsilon(\mu)_{\text{Muon} \mid \text{TrackRECO}} * \epsilon(\mu)_{\text{POG Loose ID} \mid \text{Muon}} \\ &* \epsilon(\mu)_{\text{Tight Displaced ID} \mid \text{POG Loose ID}} \end{aligned} \quad (5.10)$$

The SF maps for the terms $\epsilon(\mu)_{\text{Muon} \mid \text{TrackRECO}}$ and $\epsilon(\mu)_{\text{POG Loose ID} \mid \text{Muon}}$ are identical to those for the prompt muons (Equation 5.7). The last term corresponds to the efficiency of the tight displaced identification and isolation criteria of Table 5.4. The measurement of the last term relies on the same method as described for prompt muons at the end of Section 5.5.2 and is performed with the standard tag-and-probe technique on events from $Z \rightarrow \mu\mu$, inclusively in displacement⁵. The SFs are presented in Figure 5.23.

5.5.4 Jets and Missing Transverse Momentum

Besides leptons, also jets and the total missing transverse momentum (p_T^{miss}) in the event are important ingredients to the analysis; For signal, a jet from initial state radiation (ISR) is required to boost the sparticle pair and induce enough p_T^{miss} for the usage of p_T^{miss} based triggers. Correspondingly, the offline event selection also includes the criteria on the presence of jets, the scalar sum of jet energy (H_T) and p_T^{miss} . Furthermore, the selection of events in background control regions relies on jet selections as well as flavour tagging criteria.

Details on jet reconstruction and corrections are described in Section 4.5.1. PF jets are reconstructed using the anti- k_T algorithm with distance parameter $R = 0.4$. Particles from pileup (not associated to the PV) are removed via charged hadron subtraction. Each jet is required to have a transverse momentum of at least 25 GeV, be located within the tracker acceptance ($|\eta| < 2.4$), and satisfy identification requirements defined centrally by CMS. The jet closest to a loose electron or muon is removed in order to avoid double counting of jets and leptons, provided that the jet and lepton share a common PF candidate component. This selection conventionally is referred to as *jet cleaning*. Since LowPt electrons are

⁵The dependency on the displacement is relevant for the reconstruction scale-factors (which are hence parameterized in L_{xy}), but assumed to be negligible for the displaced muon ID.

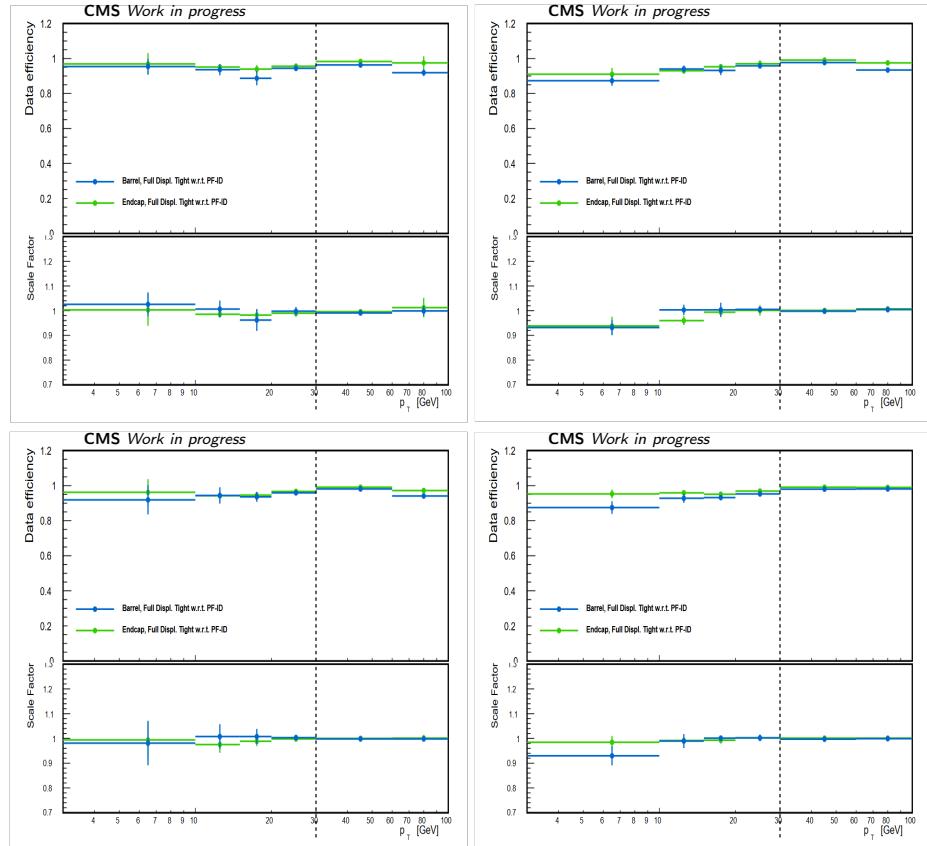


Figure 5.23: Measurements of the muon scale factors corresponding to the last term of Equation 5.7 ($\epsilon(\mu)_{\text{Tight Displaced ID} \mid \text{POG Loose ID}}$) as obtained in samples of Z boson decays, in bins of probe p_T and η . Results are shown for 2016 data periods, before (top left) and after (top right) HIP-mitigation, 2017 (bottom left), and 2018 (bottom right).

reconstructed independently from the PF algorithm, the jet cleaning with LowPt electrons consists of removing the jet closest to a loose LowPt electrons within $\Delta R < 0.4$. Jets are b tagged using the medium working point of the DeepJet discriminant, which corresponds to an efficiency of about 82% for a mistag rate in light-flavor jets of order 1%. The recommended jet energy scale corrections (derived centrally by CMS) are applied, consisting of the L1+L2L3 corrections on jets in MC simulation and observed data, and additionally L2L3 residual corrections on jets in data only.

Missing transverse momentum is reconstructed based on PUPPI candidates (Section 4.6). PUPPI-based p_T^{miss} is used (instead of PF-based p_T^{miss}), due to its improved resolution (Figure 4.12). The p_T^{miss} is Type-1 corrected by propagating the jet calibration corrections. To have a consistent treatment, the jet calibration corrections propagated to p_T^{miss} are based on AK4 PUPPI jets, while the corrections on jet related variables are based on AK4 PF jets.

5.6 Event Selection and Categorization

The targeted signal topology (Figure 5.1) consists of pair-produced electroweakinos ($\tilde{\chi}_2^0 \tilde{\chi}_1^{\pm}$) that both decay to a nearly mass-degenerate stable $\tilde{\chi}_1^0$ and an off-shell SM vector boson. As per the trigger strategy, an ISR jet in back-to-back configuration with the sparticle pair is required to induce a sufficient amount of p_T^{miss} from the undetected $\tilde{\chi}_1^0$ to use p_T^{miss} based triggers. The off-shell Z-boson is required to decay to a pair of soft opposite-sign (OS) same-flavor (SF) leptons. In the case of promptly decaying $\tilde{\chi}_2^0$, the provenance of the lepton pair is very close to the PV. Despite a reduction of multijet background that is obtained with targeting leptonic final states, there are still large SM backgrounds, which motivates the usage of tight object and event selections to gain signal sensitivity. In contrast, for long-lived $\tilde{\chi}_2^0$ scenarios, the lepton pair originates from a displaced vertex, hence the SM background rapidly decreases and its composition looks inherently different. The analysis therefore makes use of separate prompt and long-lived search categories. While the baseline selections of the search categories are similar - apart from the $\tilde{\chi}_2^0$ lifetime the topology is the same - the long-lived search category can afford to have generally looser selections due to relatively low background yields. It furthermore contains selections that exploit the secondary vertex position and lepton pair direction to decrease the background to a minimal level. The event selections for the prompt and displaced search regions are summarized in Table 5.5 and will be motivated with more detail in Sections 5.6.1 and 5.6.2, respectively.

Search for Supersymmetry with Soft Lepton Final States

As mentioned before, the invariant mass of the lepton pair is a key variable in this search, since it serves as a proxy for the mass-splitting $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ and has great discriminating power between signal and background. To maximally exploit this discrimination power and improve the robustness of the search, the analysis regions are divided into several regions of p_T^{miss} and further binned in $m_{\ell\ell}$. The p_T^{miss} binning is shown in Table 5.6 and is driven by the trigger requirements mentioned in Section 5.4.1. The plateau efficiency of the pure $p_T^{\text{miss,corr}}$ based trigger is reached at roughly $p_T^{\text{miss,corr}} = 200$ GeV, which is hence chosen as lower bound for the higher MET regions. The higher MET 2ℓ SR is further divided

Variable	Prompt 2ℓ -SR		Prompt 3ℓ -SR		Displaced 2μ -SR	
	Low-MET	Higher-MET	Low-MET	Higher-MET	Low-MET	Higher-MET
p_T^{miss} [GeV]	> 125	—	> 125	—	> 125	—
$p_T^{\text{miss,corr}}$ [GeV]	> 125	$e^{\pm}e^{\mp}/\mu^{\pm}\mu^{\mp}$	> 125	$e^{\pm}e^{\mp}\ell/\mu^{\pm}\mu^{\mp}\ell$	$\mu^{\pm}\mu^{\mp}$	$\mu^{\pm}\mu^{\mp}$
Leptons	$\mu^{\pm}\mu^{\mp}$	$e^{\pm}e^{\mp}/\mu^{\pm}\mu^{\mp}$	$\mu^{\pm}\mu^{\mp}\ell$	$e^{\pm}e^{\mp}\ell/\mu^{\pm}\mu^{\mp}\ell$	$\mu^{\pm}\mu^{\mp}$	$\mu^{\pm}\mu^{\mp}$
$p_T(\ell_1)$ [GeV] for $e(\mu)$	(5, 30)	(1(3.5), 30)	(5, 30)	(1(3.5), 30)	(5, 30)	(3, 30)
$p_T(\ell_2)$ [GeV] for $e(\mu)$	(5, 30)	(1(3.5), 30)	(5, 30)	(1(3.5), 30)	(5, 30)	(3, 30)
$p_T(\ell_3)$ [GeV] for $e(\mu)$	—	—	(5, 30)	(1(3.5), 30)	—	—
$p_T(\ell\ell)$ [GeV]	—	> 3	—	—	—	—
$\Delta R(\ell\ell)$	—	> 0.05	—	> 0.05	—	> 0.05
$M(\ell\ell)$ ($M_{\text{SFOS}}^{\text{min}}$ ($\ell\ell$) in 3ℓ) [GeV]	(4, 50)	(0.1, 50)	(4, 50)	(0.1, 50)	(4, 50)	(0.1, 50)
$M_{\text{SFAS}}^{\text{max}}(\ell\ell)$ (AS=any sign) [GeV]	veto (3, 3.2) and (9, 10.5)	—	veto (3, 3.2) and (9, 10.5)	—	veto (3, 3.2) and (9, 10.5)	—
$p_T(\ell)/M(\ell\ell)$	$p_T(\ell_2) > 0.6 + 0.25M(\ell\ell)$	$p_T(\ell_3) > 0.6 + 0.25M_{\text{SFOS}}^{\text{min}}(\ell\ell)$	—	—	—	—
$m_T(\ell, p_T^{\text{miss}})$ [GeV]	< 70	—	—	—	—	—
$\max(\Delta\phi(\ell, p_T^{\text{miss}}))$	< 1.5	—	—	—	—	—
H_T [GeV]	> 100	—	> 100	—	> 100	—
Leading jet “tight lepton veto”	✓	—	—	—	—	—
p_T^{miss}/H_T	(2/3, 1.4)	—	—	—	(2/3, 1.4)	—
$N_b(p_T > 25 \text{ GeV})$	0	0	0	0	0	—
$m_{\tau\tau}$ [GeV]	veto (0, 160)	—	—	—	—	—
Hard lepton veto ($p_T > 30 \text{ GeV}$)	—	—	—	—	✓	—
Collinearity	—	—	—	—	> -0.5	—
$\log_{10}(\Delta xy/\Delta z)$	—	—	—	—	> -1.5	—

Table 5.5: List of all event selection criteria for the prompt and displaced search categories. The label “higher-MET” collectively refers to the medium-MET, high-MET and ultra-MET bins for the prompt 2ℓ -SR, while for the prompt 3ℓ -SR and displaced 2μ -SR it refers to the single high-MET bin as defined in Table 5.6.

Region	Low-MET		Medium-MET		High-MET		Ultra-MET	
	p_T^{miss}	$p_T^{\text{miss,corr}}$						
Prompt 2ℓ -SR	> 125	(125, 200]	(200, 240]	(240, 290]	—	—	—	> 290
Prompt 3ℓ -SR	> 125	(125, 200]	—	—	> 200	—	—	—
Displaced 2μ -SR	> 125	(125, 200]	—	—	> 200	—	—	—
CRs / VRs	> 125	(125, 200]	—	—	> 200	—	—	—

Table 5.6: Definition of the MET bins in the analysis regions. The low-MET bin is not defined for the electron channel of the prompt 2ℓ -SR, the displaced 2μ -SR in 2017, the same-sign regions and the displaced opposite-sign VR for reasons related to the $\mu\mu + p_T^{\text{miss}}$ based trigger.

into a medium-, high- and ultra-MET bin to maximize sensitivity. The $\mu\mu + p_T^{\text{miss}}$ based trigger extends the acceptance down to $p_T^{\text{miss,corr}} = 125$ GeV. The $m_{\ell\ell}$ binning for the prompt SRs is based on the parametric binning strategy described in Section 5.2.2, while for the displaced SRs an alternative approach is used that profits also from the hypothesized $\tilde{\chi}_2^0$ lifetime. The final SR binnings are shown in Section 5.9.

5.6.1 Prompt Search Regions

The prompt signal regions are categorized by lepton multiplicity; For the dilepton final state, separate SRs are defined for electron and muon pairs, while the trilepton SR is inclusive in lepton flavor. Events in the low-MET regions are selected by the HLT with the $\mu\mu + p_T^{\text{miss}}$ based paths and, as such, the offline event selection for these regions includes the requirement of a pair of opposite-sign muons with $m_{\ell\ell} > 4$ GeV and $p_T > 5$ GeV. In the higher MET regions also electron pairs are selected. The p_T thresholds in the higher MET bins are as low as 1 (3.5) GeV for electrons (muons), and the lower bound on $m_{\ell\ell}$ is relaxed to 0.1 GeV. It should be emphasized that the electron thresholds can be this low due to the usage of the LowPt electron reconstruction. Backgrounds from J/ψ and ν resonances are vetoed with mass-windows and the angular lepton separation $\Delta R(\ell\ell)$ is required to be larger than 0.05 to reject lepton pairs from eg. internal conversions, in which the pairs are typically highly collimated, while maintaining high efficiency also for compressed signal hypotheses.

The selections on p_T , $m_{\ell\ell}$ and ΔR in the higher MET regions are all relaxed with respect to Ref. [89] in order to maximize acceptance to very low mass-splittings. To control the correspondingly larger background contributions, predominantly from processes with fake or non-prompt leptons (eg. W+jets or top decays), the lower bound on the trailing lepton p_T is slightly increased as function of $m_{\ell\ell}$. These background processes typically have a hard lepton from the W decay and a soft (fake) lepton from a jet. Consequently, the trailing lepton has low p_T , but $m_{\ell\ell}$ is relatively high, which causes such events to be rejected. On the other hand, leptons from signal typically have very similar momentum, and therefore, if the subleading lepton has low p_T , also $m_{\ell\ell}$ will be low. As such, signal is barely affected by this cut. These backgrounds are further reduced by exploiting the generally worse alignment of the leptons and p_T^{miss} .

The leading jet in the 2ℓ -SR is required to pass tight identification criteria for compatibility with ISR. Events with low hadronic activity are suppressed by

requiring $H_T > 100$ GeV and the selection $(2/3) < p_T^{\text{miss}}/H_T < 1.4$ was found to effectively reduce QCD events while retaining signal events (in which the p_T^{miss} is induced by the ISR jet). Events with b jets are vetoed to suppress $t\bar{t}$ background and the region $0 < m_{\tau\tau} < 160$ GeV is vetoed to reduce contamination from Z+jets events, where a boosted Z decays to two tau leptons. $Z \rightarrow \tau\tau$ decays are the only expected Drell-Yan contribution, since the tau leptons can further decay to two light leptons and four neutrinos that could induce enough p_T^{miss} to pass the relatively high p_T^{miss} bounds of the SRs. The computation of $m_{\tau\tau}$ is described in Ref. [227] and relies on the fact that the leptonic decays from energetic taus are approximately collinear. The magnitude of the two light lepton vectors are rescaled such that the leptonic pair balances the hadronic recoil, which yields a reasonable approximation of the tau momenta in $Z \rightarrow \tau\tau$ events. These momenta are used to build $m_{\tau\tau}$, where a negative value is set if the τ is determined to be going in the opposite direction to that of the light lepton.

5.6.2 Displaced Search Regions

By design, the displaced signal regions are similar to the prompt regions, but instead are based on the presence of a high quality displaced muon pair. The displaced category is agnostic to a potential third prompt or displaced soft lepton (from eg. the $\tilde{\chi}_1^\pm$ decay in signal), but does include a veto on prompt leptons that satisfy $p_T > 30$ GeV, $PFIso_{0.3}^{rel} < 0.5$ and $PFIso_{0.3}^{abs} < 5$ GeV. This veto serves to further suppress prompt SM decays, such as in W+jets events, that result in at least one prompt isolated lepton and a displaced muon pair reconstructed inside a jet. The p_T , ΔR and $m_{\ell\ell}$ cuts are largely the same as in the prompt category, but instead computed with the muon tracks constrained to the refitted secondary vertex.

A lower bound is placed on the 3D dimuon collinearity, which is defined as the cosine of the angle between the dimuon momentum vector and the vector pointing from the PV to the SV. For signal these are generally well aligned, while for background this is not always the case, in particular when the dimuon object consists of uncorrelated muons (not from the same parent particle). Finally, a lower bound is also placed on the ratio of the Δxy and Δz for both dimuon legs. For signal these distances (both with respect to the PV) typically have comparable size, while for muons from pileup the distance in the z-direction tends to be larger.

5.7 Background Estimation

5.7.1 Overview

The SM background present in the prompt and/or displaced search regions can be broadly split into three categories:

- **Processes with genuine isolated leptons from prompt decays.** The main contributions include dileptonic $t\bar{t}$ decays and $Z \rightarrow \tau\tau$, as well as di-boson processes. Both WW and ZZ production can lead to two leptons and two neutrinos (inducing p_T^{miss}), hence populate the prompt SR- 2ℓ . Contributions from the WZ process are the main background in the prompt SR- 3ℓ ($WZ \rightarrow 3\ell\nu$), but can also be found in the prompt SR- 2ℓ category in case the W decays to hadrons or an unidentified third lepton. There is only a small contribution from processes such as triboson production or $t\bar{t}+W/Z$, which have relatively low cross-section. Processes with prompt leptons are negligible in the displaced search category.
- **Processes with non-prompt or fake leptons.** This is the most important source of background, present in all search categories. It includes any SM process with at least one lepton in the final state that comes from semi-leptonic heavy flavor decays or a jet misidentified as lepton. This background mostly consists of $W/Z+jets$, or (associated) top production, but in the electron category this also includes electrons from conversions.
- **Processes with displaced muons.** The main contribution in the displaced category consists of non-prompt heavy flavor decays as mentioned above. The typical decay length⁶ of B mesons is - for a proper lifetime of $\simeq 1.6$ ps [155] and assuming a momentum of 1 GeV - roughly $500 \mu\text{m}$. However, for boosted B mesons this could increase up to $O(10)$ cm, hence this background may populate also the analysis regions corresponding to large displacement. A subdominant contribution consists of displaced muon pairs from pileup vertices. They could either both come from the same pileup vertex or from two separate vertices, in which case they are by coincidence geometrically compatible and can be refit to a common displaced secondary vertex.

The method used to estimate each of these processes depends on the nature of the background and its significance, and will be described in detail in the fol-

⁶Given as $L = \gamma\beta c\tau = \frac{p}{m}c\tau$, where γ and β are the relativistic Lorentz factors and c , p , m and τ denote the speed of light, momentum, mass and proper lifetime of the particle.

lowing sections. Processes with isolated leptons are typically well-modeled by MC simulation and can be estimated accordingly. For the main prompt backgrounds - being DY, $t\bar{t}$ and WZ - the estimations from MC are scaled to data in dedicated control regions (CR), which are defined to be orthogonal to the SRs and rich in the respective backgrounds. The remaining prompt backgrounds are estimated with MC simulation only and validated in orthogonal validation regions (VR). Processes with non-prompt or fake leptons are poorly simulated, hence are estimated with dedicated data-driven techniques. Details of this method are described in Section 5.7.3 and Appendix B. The data-driven estimations of these backgrounds are validated in dedicated CRs and VRs as well, for both the prompt and displaced analysis categories. An overview of the selection criteria of all CRs and VRs with respect to the SRs is shown in Table 5.7. Similar to the SR, the CRs and VRs are binned in $m_{\ell\ell}$ and p_T^{miss} (Table 5.6) unless indicated otherwise.

Analysis category	Region	Modified selection criteria
Prompt 2ℓ	DY-CR	$0 < m_{\tau\tau} < 160 \text{ GeV}$ No upper requirement on the lepton p_T
	$t\bar{t}$ -CR	At least one b tagged jet with $p_T > 25 \text{ GeV}$ No requirement on $m_T(\ell, p_T^{\text{miss}})$ No upper requirement on the lepton p_T
	VV-VR	$m_T(\ell, p_T^{\text{miss}}) > 90 \text{ GeV}$ $p_T(\ell_1) > 30 \text{ GeV}$ Tighter b jet veto
Prompt 3ℓ	SS-CR (high-MET only)	Same-sign requirement on lepton electric charge No requirement on $m_T(\ell, p_T^{\text{miss}})$ No requirement on $p_T(\ell)/m_{\ell\ell}$ No requirement on $\max(\Delta\phi(\ell, p_T^{\text{miss}}))$
	WZ-CR	No $M_{\text{SFOS}}^{\text{min}}(\ell)$ upper requirement at 50 GeV No $M_{\text{SEAS}}^{\text{min}}(\ell)$ requirement $p_T(\ell_1) > 30$ (37) GeV for muons (electrons) $p_T(\ell_2) > 3.5(1)$ GeV for muons (electrons) also in low-MET bin $p_T(\ell_3) > 3.5(1)$ GeV for muons (electrons) also in low-MET bin
	SS-VR (high-MET only)	Same-sign requirement on muon electric charge
Displaced 2μ	OS-VR (high-MET only)	No requirement on $\log_{10}(\Delta x/\Delta z)$
		$p_T^{\text{miss}}/H_T < (2/3)$ or $p_T^{\text{miss}}/H_T > 1.4$ or collinearity < -0.5 or at least 1 hard lepton $m_{\ell\ell} < 4 \text{ GeV}$

Table 5.7: Summary of modifications to the SR selection criteria for the prompt and displaced control and validation regions.

5.7.2 Prompt Lepton Processes

Drell-Yan

Due to the relatively high requirement on p_T^{miss} ($> 125 \text{ GeV}$), contributions from DY are only expected from $Z \rightarrow \tau\tau$ decays, where the neutrinos from subsequent tau decays induce enough p_T^{miss} . This process is reduced by requirements on the di-tau mass $m_{\tau\tau}$, but still constitutes an important background in the prompt 2ℓ -

SR. A CR rich in DY events is defined by inverting the $m_{\tau\tau}$ selection and removing the upper bound on lepton p_T to increase the statistical precision. Given the high purity of DY events, the $m_{\ell\ell}$ templates of this DY-CR are included in the final likelihood fit with the normalization of the DY MC templates left freely floating. This way the estimation of the DY background relies on MC simulated events to govern the shape of kinematic distributions, while the overall normalization is extracted from data events. The (post-fit) $m_{\ell\ell}$ templates are shown in Figure 5.24.

$t\bar{t}$ production

Production of $t\bar{t}$ with leptonic W decays constitute an important background in the prompt 2ℓ -SR as well, despite effective suppression from the b jet veto and $m_T(\ell, p_T^{\text{miss}}) < 70$ GeV requirements. Similar to the DY estimation, the $t\bar{t}$ contribution is estimated with MC simulated events and scaled to data in a high-purity CR. The $t\bar{t}$ -CR is defined by requiring at least one b tagged jet and removing the upper bound on $m_T(\ell, p_T^{\text{miss}})$. The corresponding (post-fit) $m_{\ell\ell}$ templates are shown in Figure 5.25.

Diboson production

Diboson processes populate both the prompt 2ℓ -SR and 3ℓ -SR via a variety of production and decay modes. The dominant contribution is from $WZ \rightarrow 3\ell\nu$, which is the main background in the 3ℓ -SR, but is also found in the 2ℓ -SR in case the lepton from W is a hadronically decaying tau or is not identified. This background is estimated from MC simulated events and scaled to data in a high-purity CR as well. The WZ -CR is defined similar to the 3ℓ -SR but without the upper bound on $M_{\text{SFOS}}^{\text{min}}(\ell\ell)$ to include also lepton pairs from on-shell Z decays, and orthogonality is imposed by requiring a muon (electron) with $p_T > 30$ (37) GeV. It should be noted that events in the low-MET region are triggered with $\mu\mu + p_T^{\text{miss}}$ based HLT paths that also impose upper bounds on the invariant mass (< 60 GeV). To still include events from the Z resonance for the assessment of the WZ modeling, Ref. [89] relied on double-muon triggers for the low-MET WZ -CR instead. However, these in turn imposed a p_T threshold on the leading (sub-leading) muon of 17 (8) GeV. To include more events with low p_T muons, while still accepting events from the Z resonance, events in the low-MET WZ -CR for the present analysis are instead selected with single-lepton triggers with p_T thresholds matching those of

Search for Supersymmetry with Soft Lepton Final States

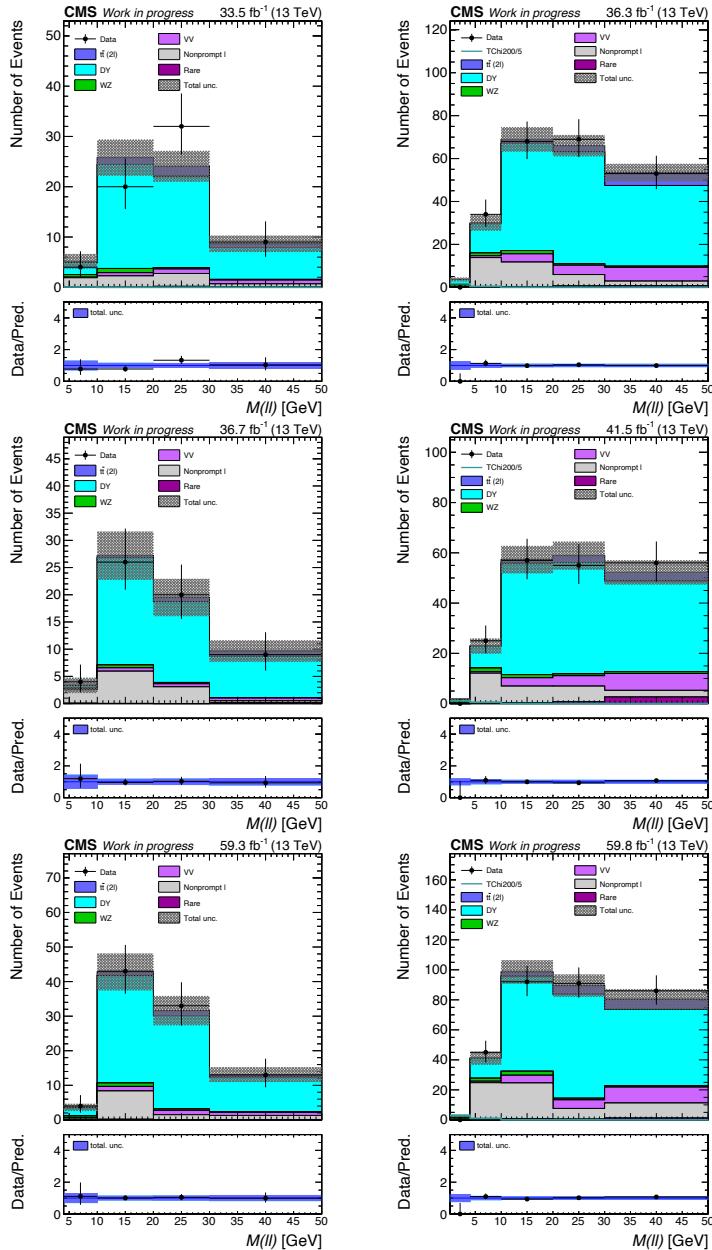


Figure 5.24: The post-fit $m_{\ell\ell}$ templates in the DY-CR. Left: Low-MET bin. Right: High-MET bin. From top to bottom: 2016, 2017, 2018 data-taking periods. Uncertainties include both the statistical and systematic components.

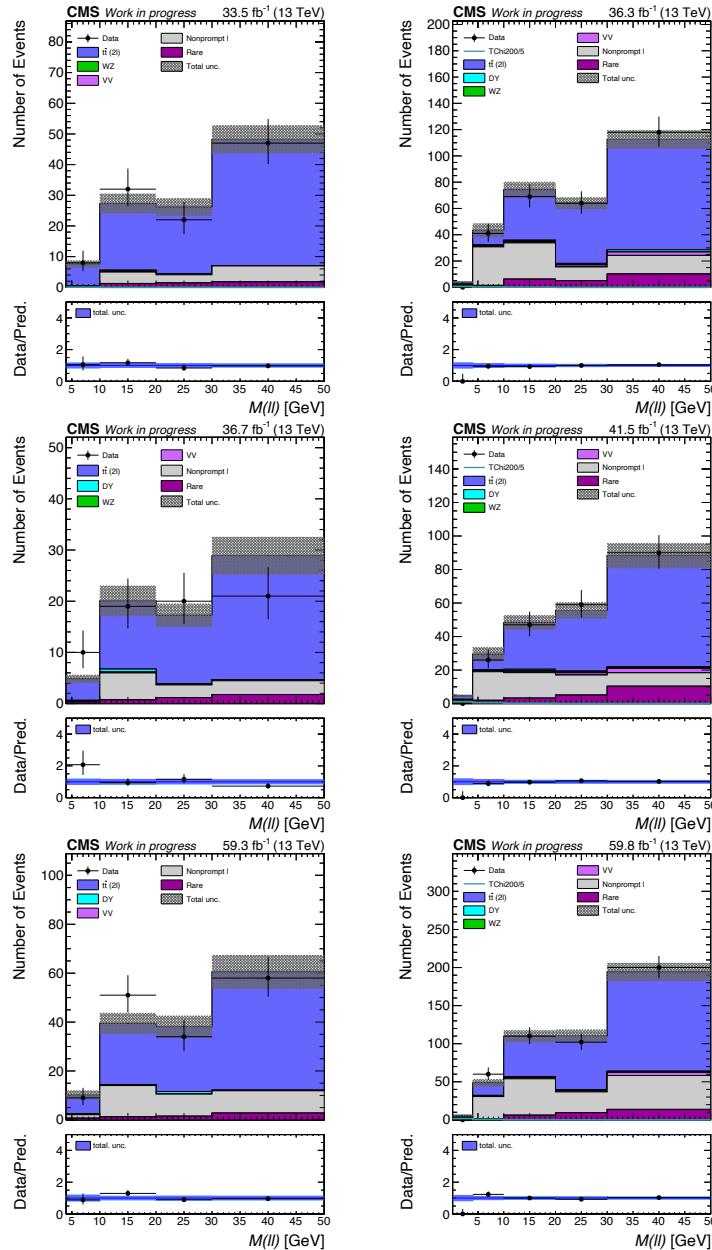


Figure 5.25: The post-fit $m_{\ell\ell}$ templates in the $t\bar{t}$ -CR. Left: Low-MET bin. Right: High-MET bin. From top to bottom: 2016, 2017, 2018 data-taking periods. Uncertainties include both the statistical and systematic components.

Search for Supersymmetry with Soft Lepton Final States

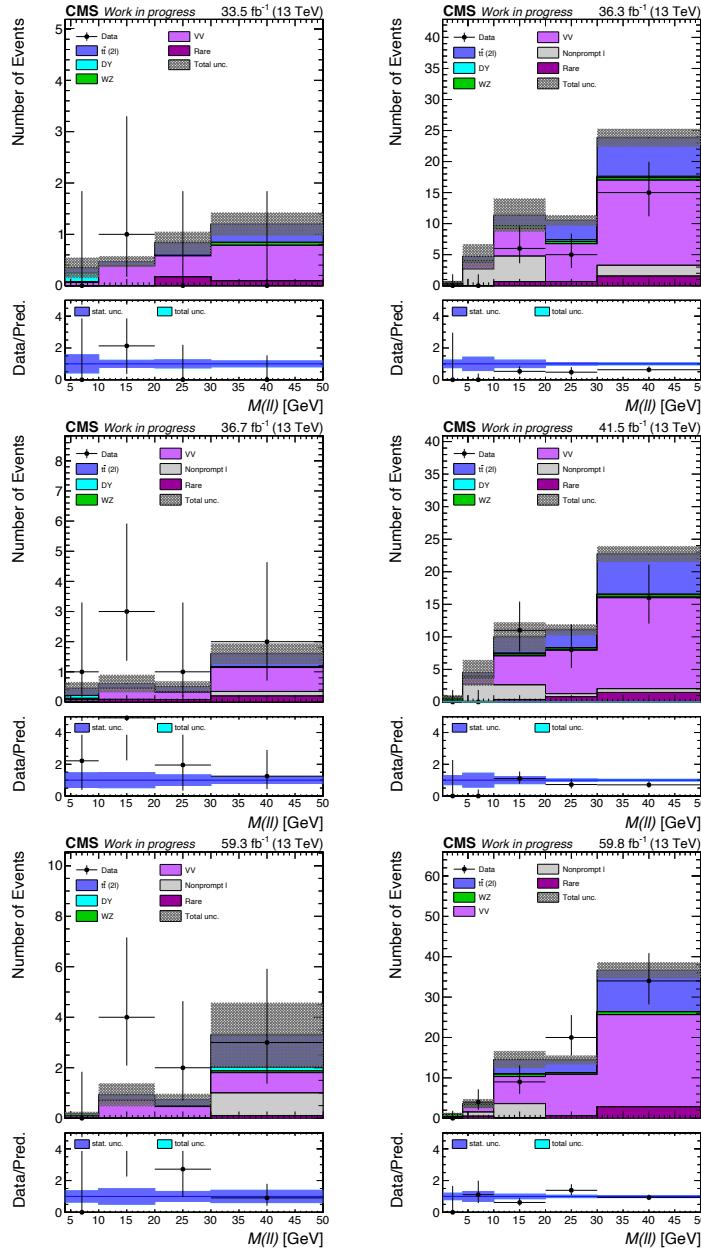


Figure 5.26: The m_{ll} templates in the VV-VR. The VV-VR is not included in the likelihood fit, but instead used to assess the MC modeling of the VV background. Left: Low-MET bin. Right: High-MET bin. From top to bottom: 2016, 2017, 2018 data-taking periods. Uncertainties include the statistical component only.

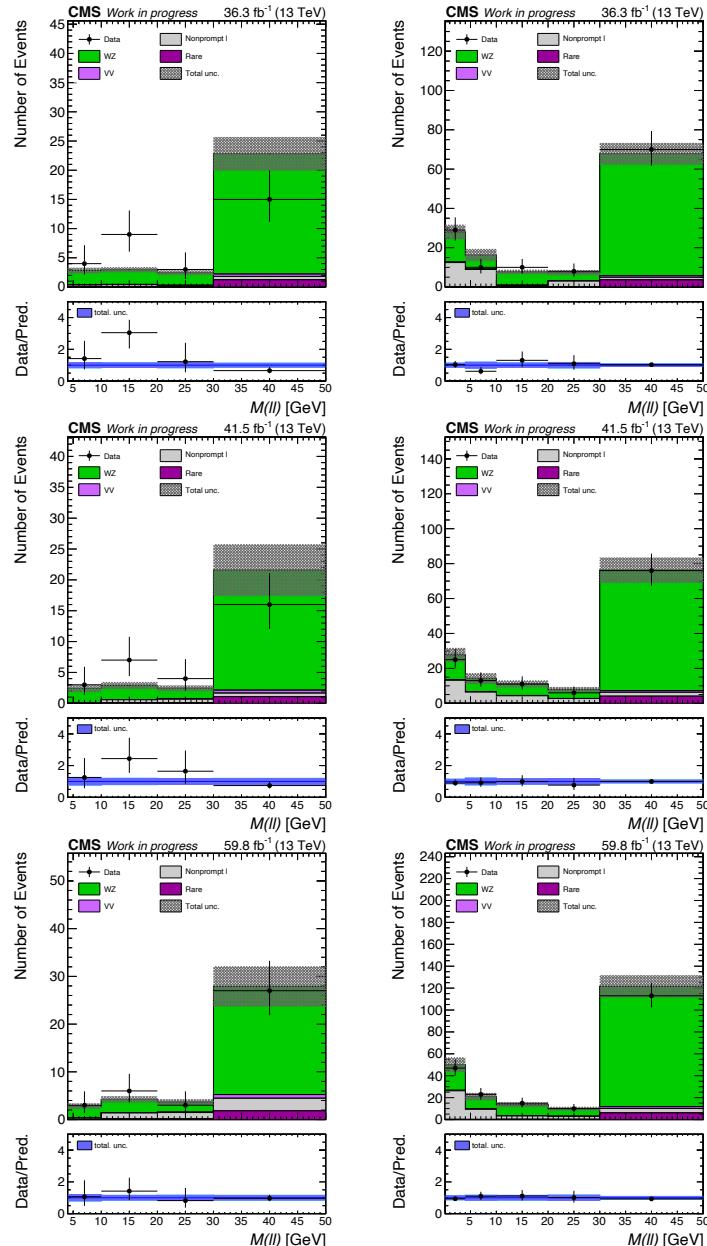


Figure 5.27: The post-fit $M_{\text{SFOS}}^{\text{min}}(\ell\ell)$ templates in the WZ-CR. Left: Low-MET bin. Right: High-MET bin. From top to bottom: 2016, 2017, 2018 data-taking periods. Uncertainties include both the statistical and systematic components.

the leading lepton (ie. 30 and 37 GeV for muons and electrons, respectively)⁷. As such, the subleading and trailing leptons can still be very soft and there is no selection on the invariant mass imposed by the trigger. This means that the full $M_{\text{SFOS}}^{\text{min}}(\ell\ell)$ spectrum can be probed also for the low-MET part of the WZ-CR and there are more events with soft muons to assess the modeling at low $M_{\text{SFOS}}^{\text{min}}(\ell\ell)$. For the sake of consistency and probing highly similar events to those populating the SR, the low-MET WZ-CR maintains the same lower bound on $M_{\text{SFOS}}^{\text{min}}(\ell\ell)$ of 4 GeV. The (post-fit) $M_{\text{SFOS}}^{\text{min}}(\ell\ell)$ templates are shown in Figure 5.27. It may be noted that the $M_{\text{SFOS}}^{\text{min}}(\ell\ell)$ modeling is still assessed down to 0.1 GeV in the high-MET part, where larger contributions are expected from $W\gamma^*$. A good agreement between the background expectation and observed data is seen.

Other diboson contributions include $WW \rightarrow 2\ell 2\nu$ and ZZ (or $Z\gamma^*$) $\rightarrow 2\ell 2\nu$. Contributions from $ZZ \rightarrow 4\ell$ or hadronically decaying bosons are smaller, since those final states either do not include neutrinos to induce enough p_T^{miss} or do not contain enough charged leptons. To validate the modeling of these backgrounds with MC simulation, a dedicated 2ℓ validation region is defined by requiring $m_T(\ell, p_T^{\text{miss}}) > 90$ GeV and inverting the leading lepton p_T threshold. The $m_{\ell\ell}$ templates for this VV-VR are shown in Figure 5.26. Despite tightening the b-jeto veto as well, the diboson purity is generally worse compared to the other CRs due to a notable contamination from $t\bar{t}$ events. Therefore, this CR is not included in the final likelihood fit, but instead is used to assess the VV modeling with MC only. Based on Figure 5.26 a conservative 50% normalization uncertainty to the VV MC is assigned in the final fit.

Rare processes

Other SM processes that yield dilepton or trilepton signatures are tW , $t\bar{t}V$, $t\bar{t}H$, tZq , tWZ and VVV processes. These processes play a minor role in the background composition of the search regions. All of these events are therefore labeled as *rare*. This background component is estimated from MC simulation and a 50% systematic uncertainty is assigned to its normalization.

5.7.3 Non-prompt or Fake Leptons

Processes with non-prompt or fake leptons are the main source of background in the analysis and are present in all search categories. It is also the type of back-

⁷In fact, the offline p_T thresholds were chosen to correspond to the plateau efficiency of the lowest-threshold unprescaled single-lepton trigger.

ground that is the most challenging to estimate for reasons described below. The term “non-prompt lepton” refers to a genuine lepton produced in semi-leptonic heavy flavor decays or in-flight decays of mesons (eg. charged kaons or pions), while the term “fake lepton” refers to any non-genuine lepton (eg. a jet) that is misidentified. Both types are obviously closely related and the terms “non-prompt” and “fake” are henceforth used interchangeably. Non-prompt and fake leptons are both associated to hadronic activity, which makes them hard to model with MC simulation. The estimation of this background therefore relies on dedicated data-driven techniques. Given the complexity of this topic, the remainder of this section provides a description of the main aspects of the estimation methods, while further details are provided in Appendix B.

Tight-to-loose Method

The baseline method for the fake lepton background estimation is often referred to as the *tight-to-loose* method [228]. The central ingredient in this method is the lepton mis-identification probability, also known as the *fake rate*, which is defined as the probability of a fake lepton that passes loose object requirements to also satisfy the tight object requirements of the analysis (Tables 5.2–5.4). The loose object requirements are typically defined by relaxing isolation and identification criteria with respect to the tight object requirements, in order to enrich the loose object collection with the targeted fake leptons. For each search region in the analysis a separate region is then defined with identical event selections, but with at least one of the selected leptons failing (any of) the tight object requirements and passing only the loose object requirements. This so-called *application region* (AR) is therefore by construction rich in fake lepton background. The fake lepton background contribution in the search regions is then obtained from the number of observed data events in the application region, scaled by a transfer factor that is a function of the fake rate. To account for the - typically small - prompt⁸ lepton contamination in the AR, the fully data-driven estimation also includes a contribution from observed data events in the SR (ie. events with all leptons passing the tight object requirements), scaled with a weight factor as well.

⁸The term “prompt” here collectively refers to genuine, isolated leptons from W , Z/γ^* or H boson decays, including leptons from the targeted signal process. Note that for long-lived signal scenarios, the term does not necessarily correspond to leptons with provenance from the PV.

Fake Rates

The fake rate itself is measured in an orthogonal fakes-enriched *measurement region* (MR), separately for prompt electrons, prompt muons and displaced muon pairs. An important feature of the fake rate is that it depends on properties of the fake lepton and - by inference - on the flavor of the mother-parton producing the fake lepton. To avoid non-closure in the fake lepton estimation due to kinematic dependencies or potentially different background compositions in the MR, AR and SR, the fake rate is parametrized as function of the main kinematic variables (most importantly p_T), and the (loose and tight) object definition are tuned to minimize the flavor dependency of the fake rate.

It should furthermore be noted that the analysis allows leptons to be very closeby ($\Delta R(\ell\ell) > 0.05$). Consequently, leptons with $\Delta R(\ell\ell) < 0.3$ enter each-others isolation cone and in most cases they are found to be decay products of the same mother particle, typically B-hadrons (light mesons) for muons (electrons). In these cases, the individual lepton selection probabilities (and thereby fake rates) are manifestly correlated, in contrast to leptons with $\Delta R(\ell\ell) > 0.3$, which can be treated as independent to an excellent approximation. This also applies to displaced muon pairs, for which the assumption of a common ancestor is already imposed by refitting them to a common vertex. For these cases, a di-object fake rate is measured that treats the lepton pairs as a composite object and is parametrized as function of kinematic variables of the pair. The measurements of the fake rates are described in Section B.1. For prompt leptons, the fake rate is parameterized as function of the lepton p_T and η , while for displaced pairs it is parameterized as function of the pair p_T and vertex displacement.

To account for the contamination of prompt lepton (or signal) processes in the observed data events used for the fake lepton estimation, the transfer factor furthermore depends on the *prompt rate*. Analogously to the fake rate, the prompt rate is defined as the probability of a given prompt object that passes the loose selection criteria to also pass the tight selection criteria. The prompt rate is similarly parametrized as function of kinematic properties of the lepton (pair), and is measured with MC simulated signal events. The prompt rates typically amount to 0.8 and 0.85 for electrons and muons, respectively (see Section B.1).

Transfer factors

As mentioned above, the transfer factors depend on both the fake rate f and prompt rate ϵ , and facilitate the estimation of the fake lepton background in the

SR based on the observed data events in the combined SR and AR. Details on the derivation of the transfer factors are given in Section B.2. In the simplest case of events with two *correlated* leptons - leptons with $(\Delta R(\ell\ell) < 0.3)$ or a displaced muon pair - the transfer factors are obtained by solving a 2×2 matrix equation and are given as:

$$W_T = \frac{(\epsilon - 1)f}{\epsilon - f} \quad (5.11)$$

$$W_L = \frac{\epsilon f}{\epsilon - f} \quad (5.12)$$

where W_T (W_L) is to be applied to events with a dilepton passing (failing) the tight object selections, ie. corresponding to SR (AR) events. For events with two *uncorrelated* leptons, or a correlated lepton pair plus a third independent lepton (which can occur in the prompt 3ℓ region), they are obtained by solving a 4×4 matrix equation:

$$W_{T,T} = \frac{(1 - \epsilon_1)(1 - \epsilon_2)f_1 f_2}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} - \frac{\epsilon_1 f_2(1 - f_1)(1 - \epsilon_2)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} - \frac{f_1 \epsilon_2(1 - \epsilon_1)(1 - f_2)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \quad (5.13)$$

$$W_{T,L} = \frac{\epsilon_1 f_2 \epsilon_2 (1 - f_1)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \quad (5.14)$$

$$W_{L,T} = \frac{\epsilon_1 f_1 \epsilon_2 (1 - f_2)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \quad (5.15)$$

$$W_{L,L} = \frac{-\epsilon_1 \epsilon_2 f_1 f_2}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \quad (5.16)$$

Finally for events in the prompt 3ℓ region with all leptons *independent*, the transfer factors are obtained analogously from a 9×9 matrix, but are not shown here given the lengthy permutations.

Semi-data-driven Variations

The object rates and transfer factors described above are the only ingredients required for the data-driven fake lepton background estimation. However, an important aspect of this search is the fact that the amount of events in the analysis regions is very low due to tight object and event selections. As a result, it may hap-

pen that the fake lepton background prediction from a fully data-driven method in a given SR bin becomes negative and statistical uncertainties are poorly defined⁹. It was found that the higher-MET prompt 2ℓ and the prompt 3ℓ regions are most susceptible to this scenario. In order to properly predict the fake lepton background in these regions, the estimation is obtained by applying the transfer factors *not* on the observed data events in the combined AR and SR, but instead on MC simulated fake lepton background events in the AR only. This variation is referred to as the *semi-data-driven* method and profits from the larger number of MC events entering the region than observed data.

To still rely on observed data to govern the shape and/or normalization of the $m_{\ell\ell}$ templates in the AR, the fake lepton MC in the AR is scaled to the observed data before applying the transfer factor. Figure 5.28 shows the higher-MET prompt 2ℓ and prompt 3ℓ AR corresponding to the 2018 data-taking period. In the 2ℓ regions the $m_{\ell\ell}$ templates from simulation poorly describe the observed data and are hence scaled bin-by-bin, in contrast to the 3ℓ regions, where the decent shape agreement permits scaling the overall normalization only (ie. inclusive in $m_{\ell\ell}$). More details on the scaling are described in Section B.3.

Since the observed data is still used to govern the shape and/or normalization of the $m_{\ell\ell}$ templates in the AR, the statistical uncertainty of the background prediction is still given by the Poisson uncertainty on the number of observed data events. However, MC simulation is used to better sample the AR phase-space, leading to more stable, positive definite predictions¹⁰. This is the crucial benefit of the semi-data-driven variation of the tight-to-loose method.

Lastly, in order to use only fake lepton MC events in the AR for the prediction, the method must be decoupled from the SR yields. As shown in Section B.3, this is done with the assumption $\epsilon \gg f$, which simplifies the transfer factors of Equations 5.11-5.16 to:

$$W_T = W_{T,T} = 0 \quad (5.17)$$

$$W_L = \frac{1}{1-f} \quad (5.18)$$

⁹Technically this is a generic feature of the tight-to-loose method resulting from negative transfer factors, such as Equation 5.16.

¹⁰The sampling of the AR phase-space with MC simulation *might* introduce a bias with respect to the observed data. However, after the MC-to-data scaling of the $m_{\ell\ell}$ templates and within the level of statistical precision available, no such bias was found in the predicted and observed lepton kinematics in the AR.

$$W_{T,L} = \frac{f_2}{1 - f_2} \quad (5.19)$$

$$W_{L,T} = \frac{f_1}{1 - f_1} \quad (5.20)$$

$$W_{L,L} = \frac{-f_1 f_2}{(1 - f_1)(1 - f_2)} \quad (5.21)$$

where it may be noted that indeed only contributions from AR events are scaled by a non-zero transfer factor.

Validating the Estimation

Finally, the fake lepton background estimation is carefully validated with samples of both observed data and MC simulated events.

The first validation is a so-called *closure test*, which consists of comparing two MC-based estimations of the SR fake lepton background; One simply corresponds to the direct estimation with MC simulation, while the other consists of applying the transfer factors, calculated from the fake rates that are measured in simulation, to the simulated fake lepton events of the AR. This test serves to extract a systematic uncertainty associated to the tight-to-loose method itself, based on the residual non-closure that is due to unparametrized kinematic dependencies of the fake rate, or different background compositions in the MR, AR and SR. Example distributions for the closure test on prompt electrons, prompt muons and displaced dimuons are shown in Figure 5.29. Based on the level of non-closure, a 40% (50%) systematic uncertainty is assigned to the fake background prediction for prompt (displaced) leptons.

A second validation is performed with observed data. For the prompt 2ℓ category a dedicated *same-sign* control region (SS-CR) is defined by requiring that the leptons have the same electric charge (sign), as shown in Table 5.7, separately for electrons and muons. This region is only defined for $p_T^{\text{miss}} > 200$ GeV due to the opposite-sign requirement of the low-MET trigger. By construction, the SS-CR is rich in events with *uncorrelated* fake leptons, which also constitutes the main fake lepton contribution in the SRs. Given the relatively high event yields of the SS-CR for both lepton flavors, this region is included in the final likelihood fit in order to constrain the systematic uncertainty on the background normalization with observed data. It should be noted, however, that the uncertainty constraint from data only applies to *uncorrelated* leptons ($\Delta R(\ell\ell) > 0.3$) due to the same-sign requirement in the SS-CR. The post-fit $m_{\ell\ell}$ templates of the prompt SS-CR (split by lepton flavor) are shown in Figure 5.30. In contrast, for the displaced

Search for Supersymmetry with Soft Lepton Final States

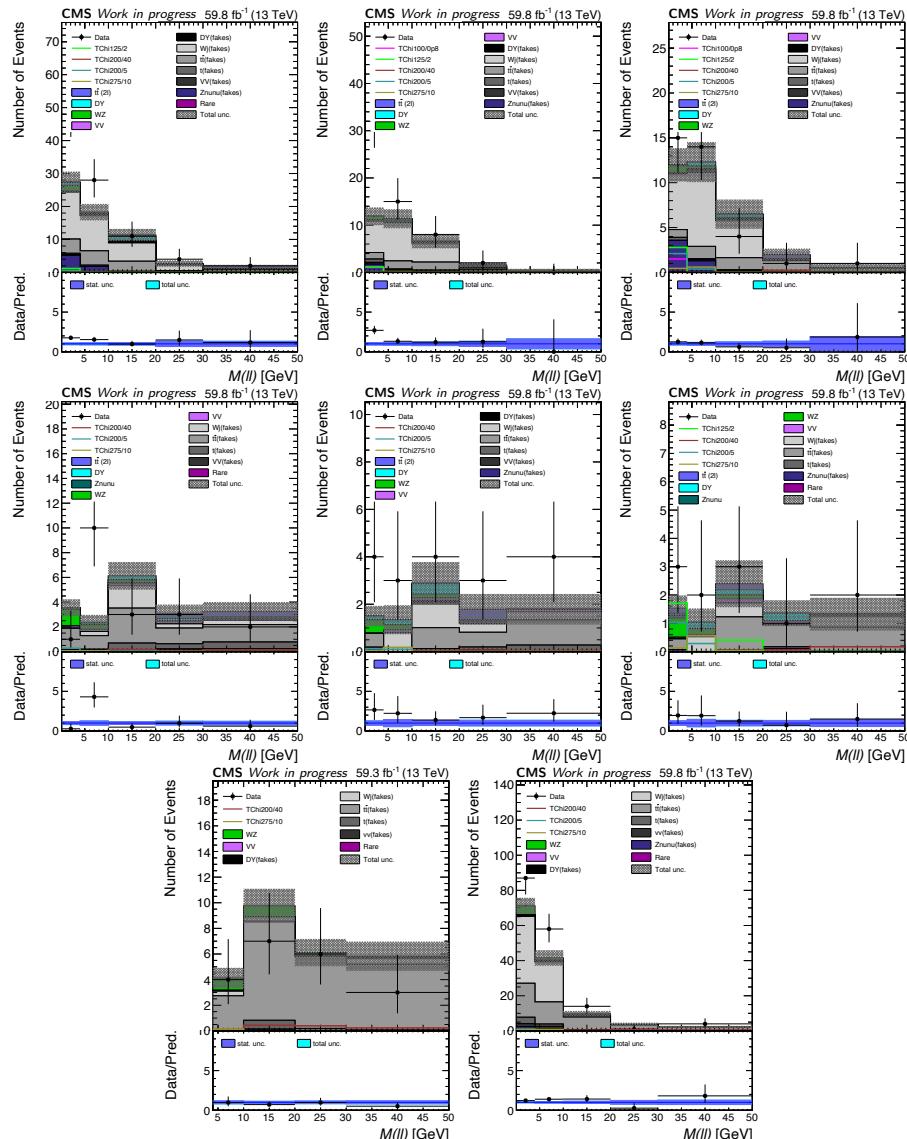


Figure 5.28: The $m_{\ell\ell}$ templates in the prompt 2ℓ -AR for electrons (top row) and muons (middle row), and in the prompt 3ℓ -AR (bottom row), corresponding to the 2018 data-taking period. From left to right the plots correspond to the medium, high and ultra-MET bins for the 2ℓ categories and the low and high-MET bins for the 3ℓ category. The fake lepton background estimates from MC are not yet scaled to the observed data. Uncertainties include the statistical component only.

2μ category the majority of the fake dimuon background contribution comes from B-meson decays, yielding *correlated* leptons. An equivalent same-sign region in the displaced category therefore contains negligible amounts of events with B-meson decays. To still validate the fake background prediction, an alternative *opposite-sign* validation region (OS-VR) is defined (Table 5.7). Given the B-meson mass, this validation region is only defined for $m_{\ell\ell} < 4$ GeV and $p_T^{\text{miss}} > 200$ GeV (the latter again due to the trigger requirements). The $m_{\ell\ell}$ templates for the displaced 2μ OS-VR are shown in Figure 5.31 (left). Overall good agreement is observed.

5.7.4 Displaced Lepton Processes

As mentioned above, the main contribution of background in the displaced search category consists of non-prompt heavy flavor decays. Depending on the boost of the (B) hadron, this contribution can populate regions with displacements up to $O(10)$ cm. The estimation of this background is performed with the tight-to-loose method, and validated with both MC simulation and observed data events, as already described in Section 5.7.3.

The only other source of background populating the displaced analysis regions is from pileup tracks. The tracks may either come from the decay of the same ancestor particle that is produced in a single pileup vertex (eg. a pileup B-meson), or be uncorrelated. In the latter case, they originate from two separate pileup vertices, but are still geometrically compatible by coincidence, such that they are refit to a common displaced secondary vertex. Muon pairs from pileup are generally softer, and can have also have the same electric charge in case they are from different pileup vertices. Although conceptually the tight-to-loose method could still be used to estimate this fake lepton contribution, the provenance from pileup gives a technical complication; As described in Section 5.5.3, the isolation sum for displaced muons is based on the standard PF isolation computation, which assumes provenance from the PV. Since charged hadrons from pileup are excluded from the isolation sum, the muons from pileup may *appear* more isolated. Consequently, the probability for a dimuon that passes the loose object requirements to also pass the tight requirements increases, yielding artificially high fake rates and making the tight-to-loose method unsuitable for the background estimate. To test the proper modeling of this background with MC simulation, a pileup-enriched validation region (SS-VR) is constructed by imposing a same-sign requirement on the muon pair, as shown in Table 5.7. Given the opposite-sign requirement on the low-MET trigger, this SS-VR is only defined for $p_T^{\text{miss}} > 200$ GeV. The

Search for Supersymmetry with Soft Lepton Final States

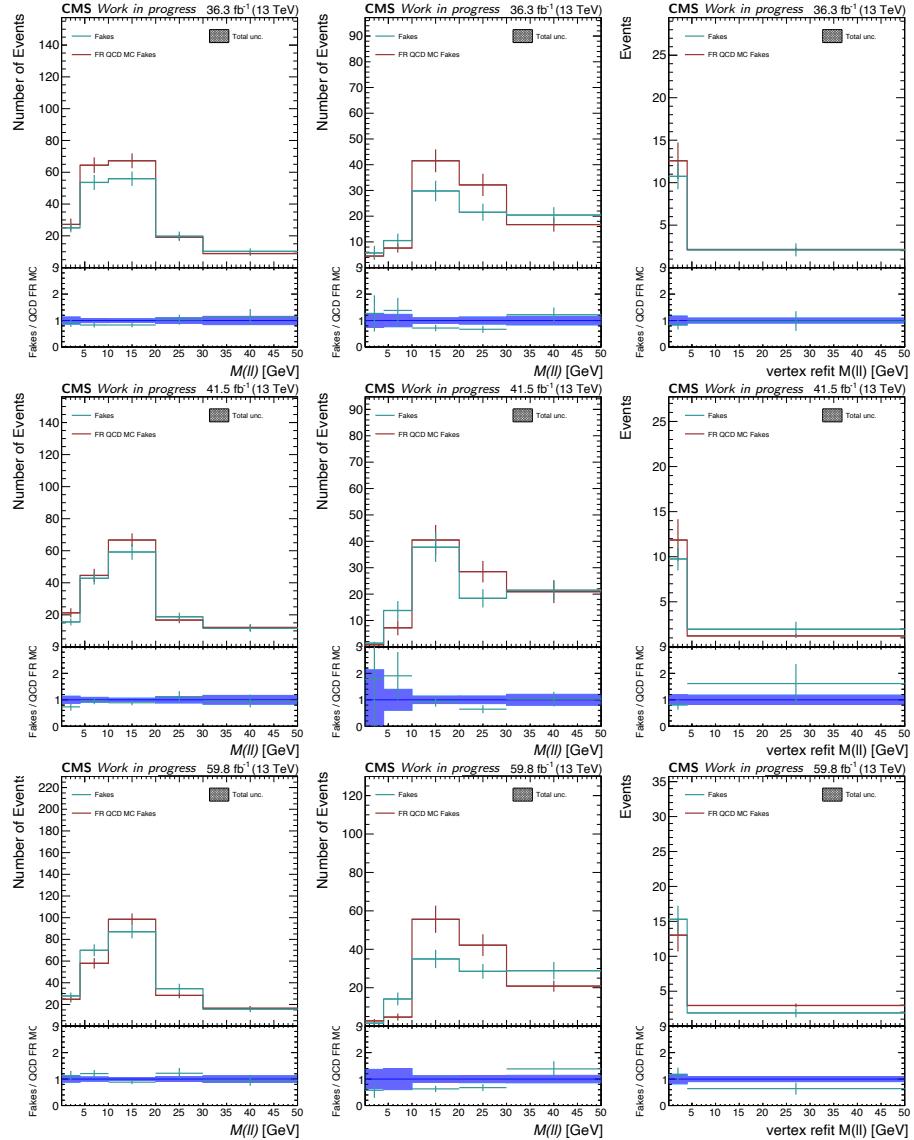


Figure 5.29: Closure tests of the tight-to-loose method, corresponding to the prompt 2ℓ -SR for electrons (left column) and muons (middle column), and the displaced 2μ -SR. From top to bottom: 2016, 2017 and 2018 data-taking periods. Uncertainties include the statistical component only.

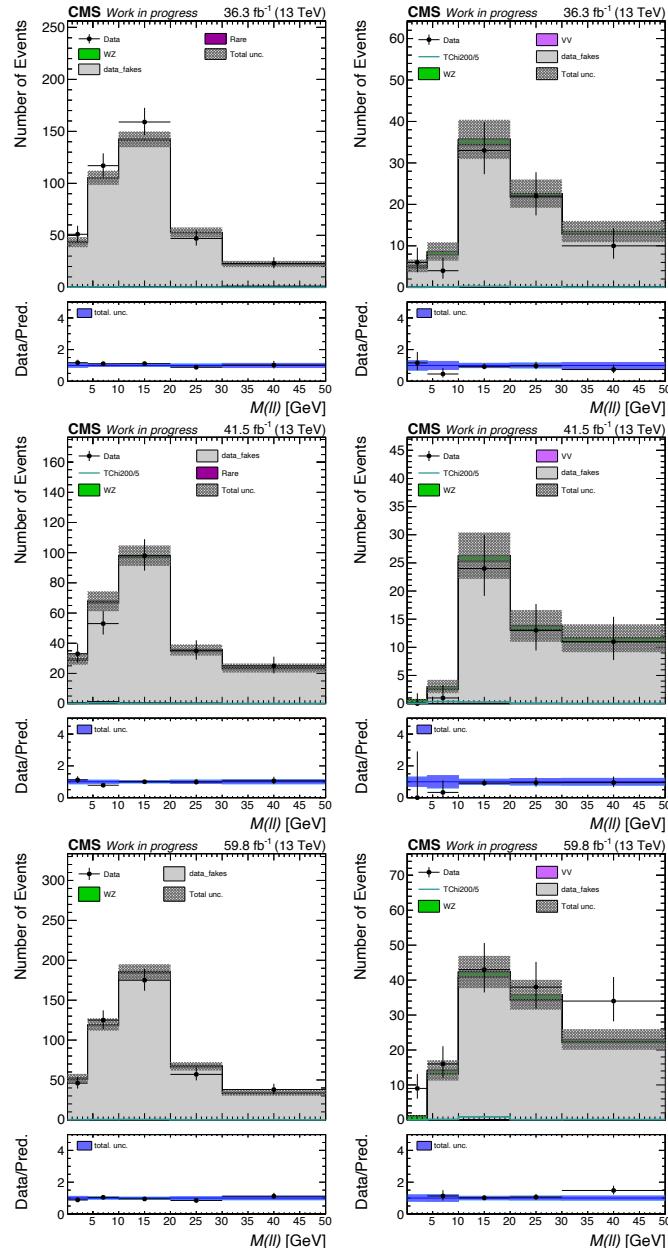


Figure 5.30: The post-fit $m_{\ell\ell}$ templates in the prompt 2ℓ SS-CR for electrons (left column) and muons (right column). From top to bottom: 2016, 2017 and 2018 data-taking periods. Uncertainties include both the statistical and systematic components.

Search for Supersymmetry with Soft Lepton Final States

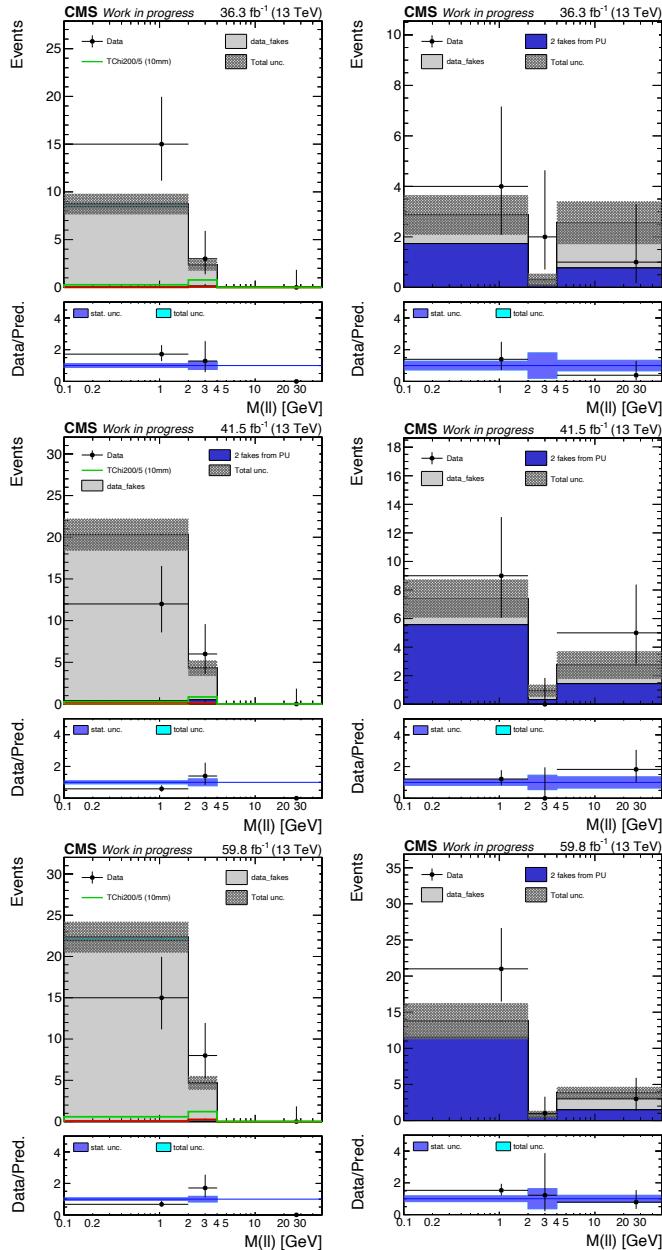


Figure 5.31: The $m_{\ell\ell}$ templates in the displaced 2μ OS-VR (left column) and SS-VR (right column), used to assess the data-driven estimate of B-meson decays and the MC estimate of pileup contributions, respectively. From top to bottom: 2016, 2017 and 2018 data-taking periods. Uncertainties include the statistical component only.

$m_{\ell\ell}$ templates are shown in Figure 5.31 (right). Given the good modeling of this background in the pileup-enriched SS-VR, this subdominant SR contribution is estimated directly from MC instead.

5.8 Systematic Uncertainties

The dominant uncertainty in the analysis is of statistical origin, coming from the limited statistical precision of the MC simulated samples and observed data events. Additional systematic uncertainties are of experimental origin, related to eg. detector effects, or are related to potential mismodeling of signal and background processes. Below follows a description of the systematic uncertainties considered in the analysis.

Trigger efficiency - In order to correct for potential mismodeling of the online event selection in simulated events, data-to-simulation scale factors are applied to all MC events, as described in Section 5.4.1. Each scale factor come with a systematic uncertainty that mainly originates from the statistical precision of the samples used for the trigger efficiency measurements. For the scale-factors associated to the p_T^{miss} based trigger path and the p_T^{miss} component of the $\mu\mu + p_T^{\text{miss}}$ based trigger path, a 2% uncertainty is assigned to the plateau efficiency region (ie. $p_T^{\text{miss,corr}} > 150$ GeV and $p_T^{\text{miss,corr}} > 250$ GeV, respectively). This is inflated to 5% at lower $p_T^{\text{miss,corr}}$ values, corresponding to the efficiency turn-on. For the muon component of the $\mu\mu + p_T^{\text{miss}}$ based trigger, a 2% uncertainty is assigned per muon leg. Lastly, for the single-lepton triggers, 2% uncertainty is applied as well. A nuisance parameter is assigned in the final likelihood fit for the combined trigger scale factor uncertainty. Since the lepton reconstruction at the HLT and the pileup profile that drives the trigger performance changed across years, a separate uncorrelated nuisance parameter is assigned for each year of data-taking.

Trigger prefire - During the 2016 and 2017 data-taking periods, a gradual shift was observed in the timing of L1 trigger primitives from the ECAL. This caused a large fraction of e/γ objects with $|\eta| > 2.5$ to be associated to the previous bunch crossing. Per-event correction factors have been derived centrally by CMS and are applied to MC simulated events for these years. The associated uncertainties are implemented as shape uncertainties on the $m_{\ell\ell}$ templates in the likelihood fit.

Lepton efficiency - Uncertainties associated to the selection efficiency scale factors for the full set of tight object requirements are already shown in Section 5.5. The uncertainties stem from the statistical precision of the samples used in the

efficiency measurements, and from the functional forms used to fit the signal and background distributions in the tag-and-probe method. The uncertainties are assigned on a per-object basis.

Jet energy corrections and b tagging efficiency - The jet energy corrections and b tagging scale factors are derived centrally by CMS based on a series of dedicated measurements (Section 4.5.1). The associated uncertainties correspond to eg. mismodeling of the pileup profile or detector effects and may depend on the jet p_T and η . All JEC-related uncertainties are grouped, considered correlated across years, and implemented as shape uncertainties on the $m_{\ell\ell}$ template in the final likelihood fit. For the b tagging of jets, uncertainties are introduced by background contamination in samples used for the tagging efficiency measurements, as well as the statistical precision of the samples. These sources are treated as independent, but correlated across years.

Pileup modeling - Weights are applied to MC simulated events to correct the distribution of the number of vertices to that measured in data. The number of interactions per bunch crossing is estimated from the total inelastic cross section, which was measured to be $\sigma_{\text{MinBias}} = 69.2$ mb, with an uncertainty of 4.6% [229]. The pileup uncertainty is estimated by propagating the minimum bias cross section uncertainty to the pileup weights and using their variation as a shape uncertainty on the $m_{\ell\ell}$ templates.

Luminosity measurement - The uncertainty on the luminosity measurements are incorporated as flat normalization uncertainties for both the background and signal predictions. The uncertainties are 1.2% for 2016 (both for pre- and post-VFP periods), 2.3% for 2017 and 2.5% and 2018.

Prompt background modeling - For each of the main prompt lepton backgrounds (DY, $t\bar{t}$ and WZ), a dedicated high-purity control region is designed. These control regions are included in the final likelihood fit, and the normalization of the MC estimations is left freely floating within a factor 2 of the prefit normalization. This allows the normalization to be governed by the observed data in the respective control regions.

Diboson and rare background modeling - A conservative 50% normalization uncertainty is assigned to the estimation of the remaining diboson contribution (mainly $WW \rightarrow 2\ell 2\nu$ and $ZZ \rightarrow 2\ell 2\nu$) as well as rare background processes.

Non-prompt or fake lepton background estimation - The non-prompt lepton background is estimated via a (semi-)data driven procedure that introduces several sources of uncertainty. The main component is related to the performance of the tight-to-loose method, which is studied in closure tests (Figure 5.29). Based

on the level of non-closure, a normalization uncertainty is assigned to the prediction of the fake lepton background, separately for prompt leptons and displaced dimuons. These amount to 40% and 50%, respectively. At the time of writing, the remaining lesser components are still to be implemented; Firstly, for the semi-data driven estimates (in the prompt 3ℓ and higher-MET prompt 2ℓ SRs), the AR fake lepton MC events are scaled to the observed data, allowing the latter to govern the normalization and/or shape of the fake lepton $m_{\ell\ell}$ templates. As such, the statistical uncertainty of the background prediction is given by the Poisson uncertainty on the number of observed data events, whereas it is currently still based on the number of simulated events. Secondly, in the prompt muon category, the scaling is performed inclusively in p_T^{miss} for robustness against statistical fluctuations. This assumes that the fake lepton $m_{\ell\ell}$ templates are independent of p_T^{miss} . To account for potential discrepancies in the $m_{\ell\ell}$ shapes as function p_T^{miss} , a shape uncertainty will be assigned based on the simulated fake lepton $m_{\ell\ell}$ templates. Lastly, to constrain the fake lepton background in the prompt analysis, the prompt same-sign control region (one for each lepton flavor) is included in the final fit in a single MET bin ($p_T^{\text{miss,corr}} > 200$ GeV). Due to the same-sign requirement, the constraint from the SS-CR should apply only to independent leptons ($\Delta R(\ell\ell) > 0.3$). However, at the time of writing it is still applied inclusively in ΔR . The full implementation of the above uncertainties is expected to have only minor effects compared to the main normalization uncertainty that is already assigned. Finally, uncertainties on the measured fake (and prompt) rates associated to limited statistical precision of the samples used in the measurements are generally small compared to the main normalization uncertainty. This uncertainty is therefore not propagated to the (semi-)data-driven estimates.

Signal modeling - The theoretical uncertainty from the choice of factorization and normalization scales is propagated to the predicted $m_{\ell\ell}$ templates by varying these scales up and down by a factor 2. The differences between the resulting $m_{\ell\ell}$ templates and the nominal prediction will be used as shape uncertainty in the likelihood fit. At the time of writing these are not yet implemented.

5.9 Results

The $m_{\ell\ell}$ distributions of prompt 2ℓ , prompt 3ℓ and displaced 2μ search regions, with the (semi-)data-driven fake lepton background estimations are shown in Figures 5.32-5.35. These distributions, together with those of the prompt 2ℓ DY-CR, $t\bar{t}$ -CR, SS-CR and prompt 3ℓ WZ-CR are used for the binned likelihood fit of the

Search for Supersymmetry with Soft Lepton Final States

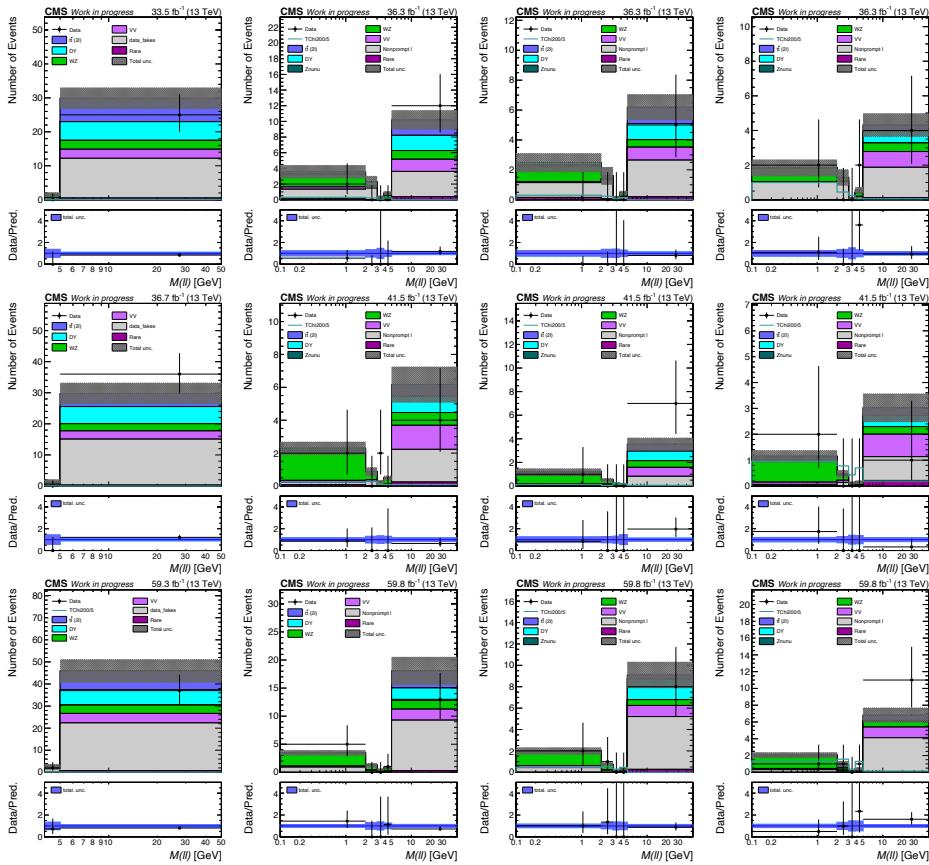


Figure 5.32: The post-fit m_{ll} distributions in the muon channel of the prompt 2ℓ -SR. From left to right: low, medium, high and ultra-MET bins. From top to bottom: 2016, 2017 and 2018 data-taking periods. The distributions are based on the parametric binnings derived for signal mass-points with $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^\pm) = 5$ GeV. The pre-fit signal distribution for $m_{\tilde{\chi}_2^0} = m_{\tilde{\chi}_1^\pm} = 200$ GeV with $m_{\tilde{\chi}_2^0} = 195$ GeV is overlaid for illustration. Uncertainties include both the statistical and systematic components. At the time of writing, minor components related to the fake lepton background and signal modeling uncertainties are still being implemented (see Section 5.8).

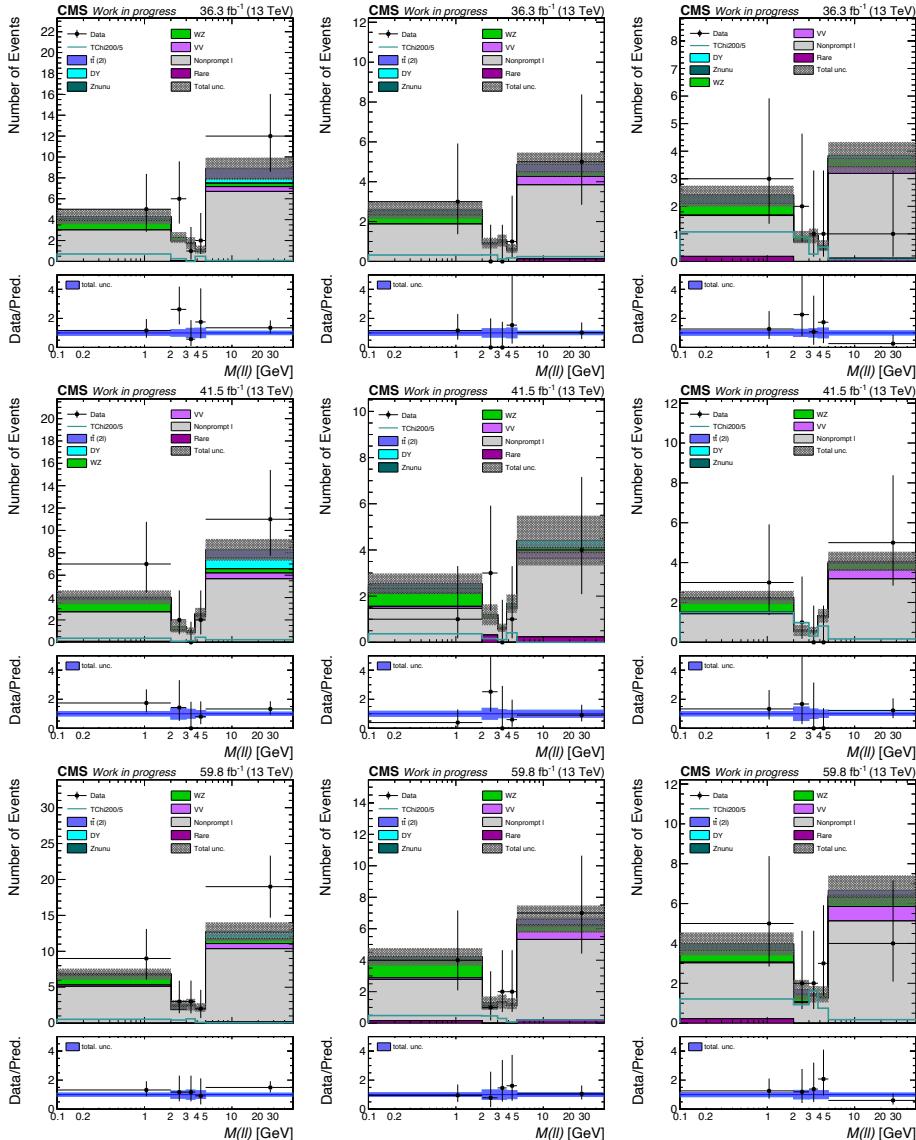


Figure 5.33: The post-fit $m_{\ell\ell}$ distributions in the electron channel of the prompt 2ℓ -SR. From left to right: medium, high and ultra-MET bins. From top to bottom: 2016, 2017 and 2018 data-taking periods. The distributions are based on the parametric binnings derived for signal mass-points with $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) = 5$ GeV. The pre-fit signal distribution for $m_{\tilde{\chi}_2^0} = m_{\tilde{\chi}_1^\pm} = 200$ GeV with $m_{\tilde{\chi}_2^0} = 195$ GeV is overlaid for illustration. Uncertainties include both the statistical and systematic components. At the time of writing, minor components related to the fake lepton background and signal modeling uncertainties are still being implemented (see Section 5.8).

Search for Supersymmetry with Soft Lepton Final States

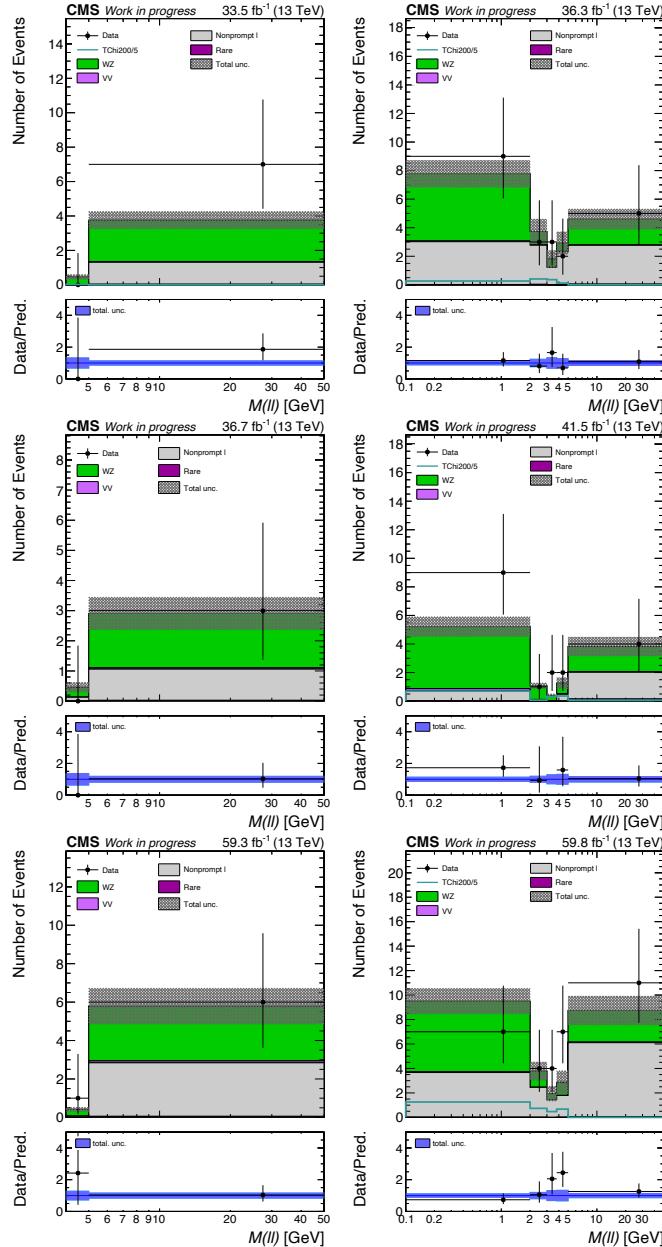


Figure 5.34: The post-fit $m_{\ell\ell}$ distributions in the prompt 3ℓ -SR, corresponding to the low-MET (left) and high-MET (right) bins. From top to bottom: 2016, 2017 and 2018 data-taking periods. The distributions are based on the parametric binnings derived for signal mass-points with $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) = 5$ GeV. The pre-fit signal distribution for $m_{\tilde{\chi}_2^0} = m_{\tilde{\chi}_1^\pm} = 200$ GeV with $m_{\tilde{\chi}_2^0} = 195$ GeV is overlaid for illustration. Uncertainties include both the statistical and systematic components. At the time of writing, minor components related to the fake lepton background and signal modeling uncertainties are still being implemented (see Section 5.8).

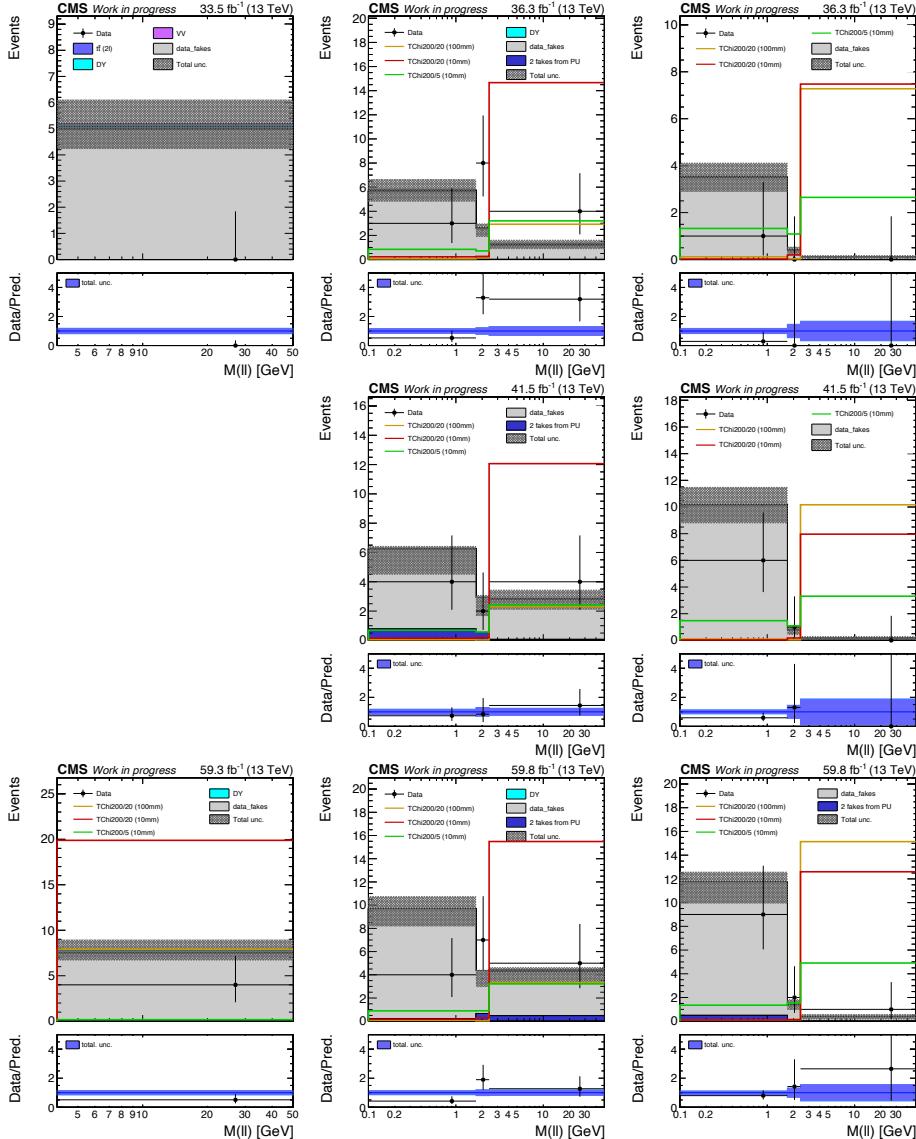


Figure 5.35: The post-fit $m_{\ell\ell}$ distributions in the displaced 2μ SR. From left to right: low-MET bin (inclusive in L_{xy}), high-MET bin with $L_{xy} < 0.9$ cm and high-MET bin with $L_{xy} > 0.9$ cm. From top to bottom: 2016, 2017 and 2018 data-taking periods. No low-MET region is defined for 2017, since the DCA-based trigger was disabled during most of the data-taking period. The distributions are based on the parametric binnings derived for signal mass-points with $c\tau \geq 10$ mm. Several pre-fit signal distributions are overlaid for illustration. Uncertainties include both the statistical and systematic components. At the time of writing, minor components related to the fake lepton background and signal modeling uncertainties are still being implemented (see Section 5.8).

signal and background expectations to the data. The uncertainties are included as nuisance parameters in the fit, as described above.

To maximize the signal sensitivity, the $m_{\ell\ell}$ binnings have been optimized independently for each signal hypothesis targeted in this search, based on the kinematic distributions of the signal and expected background in each of the search regions. For the prompt SRs, the $m_{\ell\ell}$ binnings are derived as a function of the mass-splitting, while the $m_{\ell\ell}$ binnings of the displaced SRs depend on the lifetime of the targeted signal hypothesis. Details of the SR $m_{\ell\ell}$ binning derivations are described below. The SR distributions shown in this section correspond to signal scenarios with mass-splitting $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) = 5$ GeV for the prompt analysis category and to $c\tau \geq 10$ mm for the displaced category. The $m_{\ell\ell}$ binnings of the CRs included in the fit are those shown in Section 5.7, and do not depend on the signal scenario.

5.9.1 Binning optimization

To improve the signal sensitivity over the full parameter space targeted in this search, the SR binnings are optimized independently for each signal hypothesis.

For the prompt SRs, the $m_{\ell\ell}$ binning is optimized using the exact same strategy as described in Section 5.2.2. This *parametric binning strategy* aims to better exploit the shape of the $m_{\ell\ell}$ distribution, by designing a unique $m_{\ell\ell}$ binning for each $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ hypothesis, separately for each lepton category and MET bin, while also guaranteeing a minimal level of statistical precision of the background estimates.

For the displaced SRs, the binning relies on a modified version of this strategy. Since these search regions have relatively low SM background yields (due to the requirement of displacement), and are mainly populated with backgrounds from B-decays (sharply peaking at $m_{\ell\ell} < 5$ GeV), the procedure as developed for the prompt search category is no longer optimal. Furthermore, the above strategy does not profit from the non-zero lifetime of long-lived $\tilde{\chi}_2^0$. As such, the alternative strategy to determine an optimal set of bins for the displaced search regions, based on a grid search, is defined as follows:

1. A two-dimensional grid of potential bin-boundaries is defined in the range $\Delta m = [0.1, 50]$ GeV and $L_{xy} = [0, 10]$ cm, in steps of 0.1 GeV and 0.3 cm respectively. Any combination of bin-boundaries that yields three bins in $m_{\ell\ell}$ and two bins in L_{xy} is used as candidate binning. The L_{xy} variable was chosen for the parametrization of the displacement, since it was found to

give better sensitivity overall compared to $\sigma_{L_{xy}}$.

2. Candidate binnings that have large statistical fluctuations are rejected. The same figure of merit is used as in the derivation of the prompt SR binnings, but with a lower threshold given the low expected background yields in the displaced SRs: $N_{eq} > 1.0$.
3. For all signal hypotheses considered in this search, the remaining candidate binnings are tested for their expected signal significance, defined as:

$$S_{global} = \sqrt{\sum_{i,j} \left(\frac{s_{ij}}{\sqrt{b_{ij}}} \right)^2} \quad (5.22)$$

where the sum runs over the $(3 \times 2 = 6)$ $m_{\ell\ell}$ and L_{xy} bins of the candidate binning. The optimal candidate is the candidate with highest value of S_{global} .

For the sake of robustness, in the high-MET displaced SR only two candidate binnings are selected, which represent the optimal binning for the majority of the signal hypotheses. In the low-MET displaced SR, no candidate was found that satisfies $N_{eq} > 1.0$ due to the very low expected background, so a single bin is used. As such, the final binnings for the displaced SR are:

- high-MET, $c\tau < 10$ mm: $L_{xy} = [0, 0.3, \infty)$ cm, $m_{\ell\ell} = [0.1, 2.5, 3.7, 50]$ GeV
- high-MET, $c\tau \geq 10$ mm: $L_{xy} = [0, 0.9, \infty)$ cm $m_{\ell\ell} = [0.1, 1.7, 2.4, 50]$ GeV
- low-MET: $L_{xy} = [0, \infty)$ cm, $m_{\ell\ell} = [0.1, 50]$ GeV

5.9.2 Signal Interpretations

The observed data agree with the SM background expectations within the 1σ uncertainty bands in nearly all SR bins (post-fit). The results are interpreted in terms of a simplified supersymmetric model for the production of a wino-like $\tilde{\chi}_2^0 \tilde{\chi}_1^\pm$ pair, both decaying to a bino-like $\tilde{\chi}_1^0$, as described in Section 5.1. Limits as function of the mass-splitting $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$, $\tilde{\chi}_2^0$ mass and life-time are set at 95% confidence level using the modified frequentist approach (Appendix A). At the time of writing, the MC signal samples for the higgsino scenario are still being produced, hence the higgsino interpretation will be added in future.

The expected and observed exclusion limits corresponding the assumption of promptly decaying $\tilde{\chi}_2^0$ are shown in Figure 5.36 and compared with the results of

Ref. [89] (green line). Major improvements in sensitivity for highly compressed signal scenarios ($\Delta m < 5$ GeV) are obtained due to the usage of electrons down to $p_T = 1$ GeV, optimized identification criteria and new event selections. This search now excludes mass-splittings as low as 900 MeV for a $\tilde{\chi}_2^0$ mass of 100 GeV. At higher mass-splittings ($30 < \Delta m < 50$ GeV) the observed exclusion limit is weaker than the expected one. This is due to upward data fluctuations in specific $m_{\ell\ell}$ bins of the high- and ultra-MET 2ℓ -SR (muon channel, $5 < m_{\ell\ell} < 50$ GeV), the 3ℓ -SR ($5 < m_{\ell\ell} < 50$ GeV) and the low-MET WZ-CR ($10 < m_{\ell\ell} < 20$ GeV). These regions are largely unchanged with respect to Ref. [89], in which similar fluctuations were observed.

Limits corresponding to the same simplified model, but assuming $\tilde{\chi}_2^0$ lifetimes up to $c\tau = 100$ mm, are shown in Figures 5.37 and 5.38. The new displaced search category, based on the presence of soft displaced muon pairs, provides a unique handle to effectively probe mass-compressed *and* long-lived signal scenarios. This is in shear contrast with Ref. [89], which had negligible sensitivity for signals with $c\tau \geq 1$ mm. The peak sensitivity of the present analysis is for signals with $c\tau$

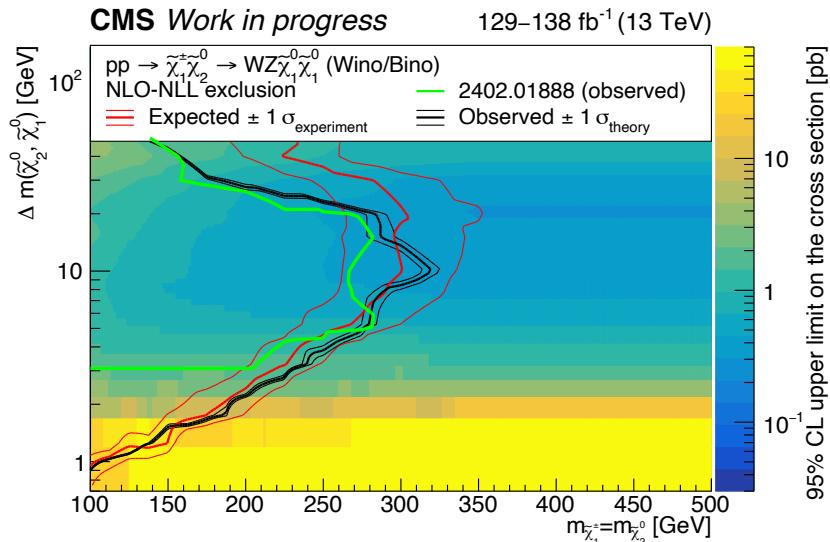


Figure 5.36: Observed and expected exclusion limits at 95% confidence level, based on the simplified model for chargino-neutralino production with $c\tau(\tilde{\chi}_2^0) = 0$ mm and assuming NLO+NLL wino production cross-sections. The solid green line corresponds to the observed exclusion limits of Ref. [89]. At the time of writing, minor components related to the fake lepton background and signal modeling uncertainties are still being implemented (see Section 5.8).

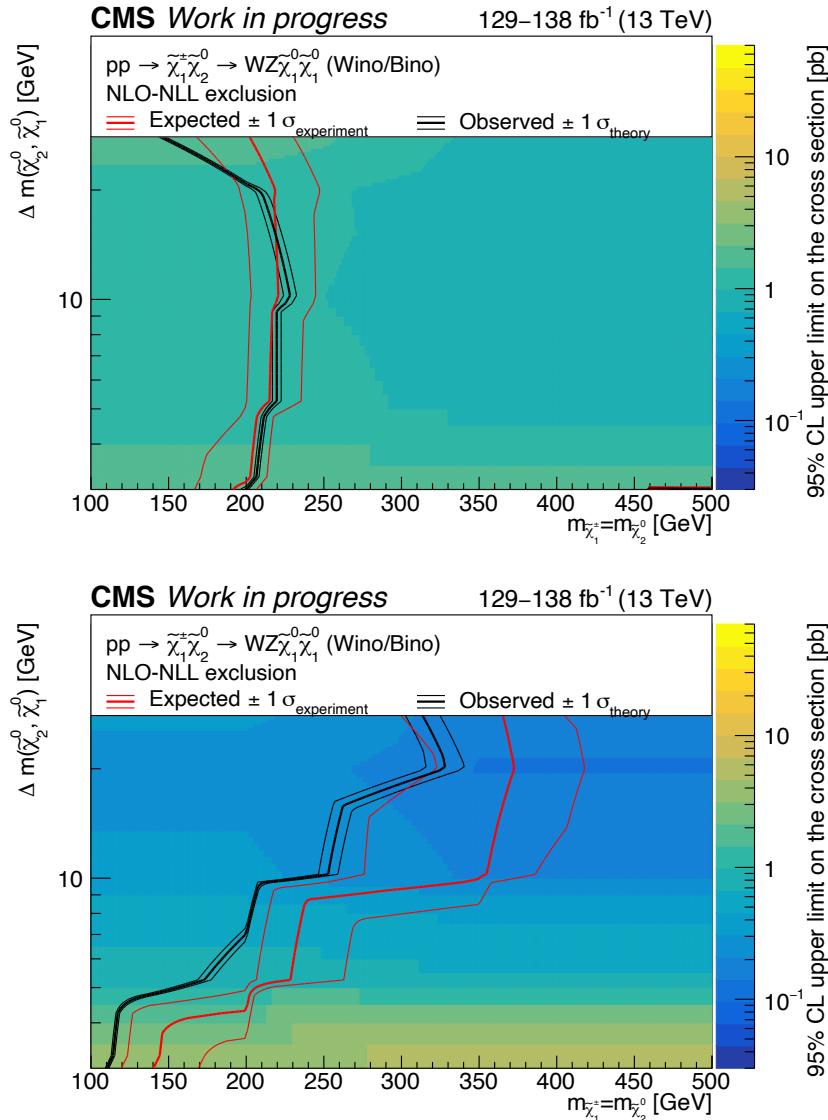


Figure 5.37: Observed and expected exclusion limits at 95% confidence level, based on the simplified model for chargino-neutralino production with $c\tau(\tilde{\chi}_2^0) = 0.1$ mm (top) and 1.0 mm (bottom), and assuming NLO+NLL wino production cross-sections. At the time of writing, minor components related to the fake lepton background and signal modeling uncertainties are still being implemented (see Section 5.8).

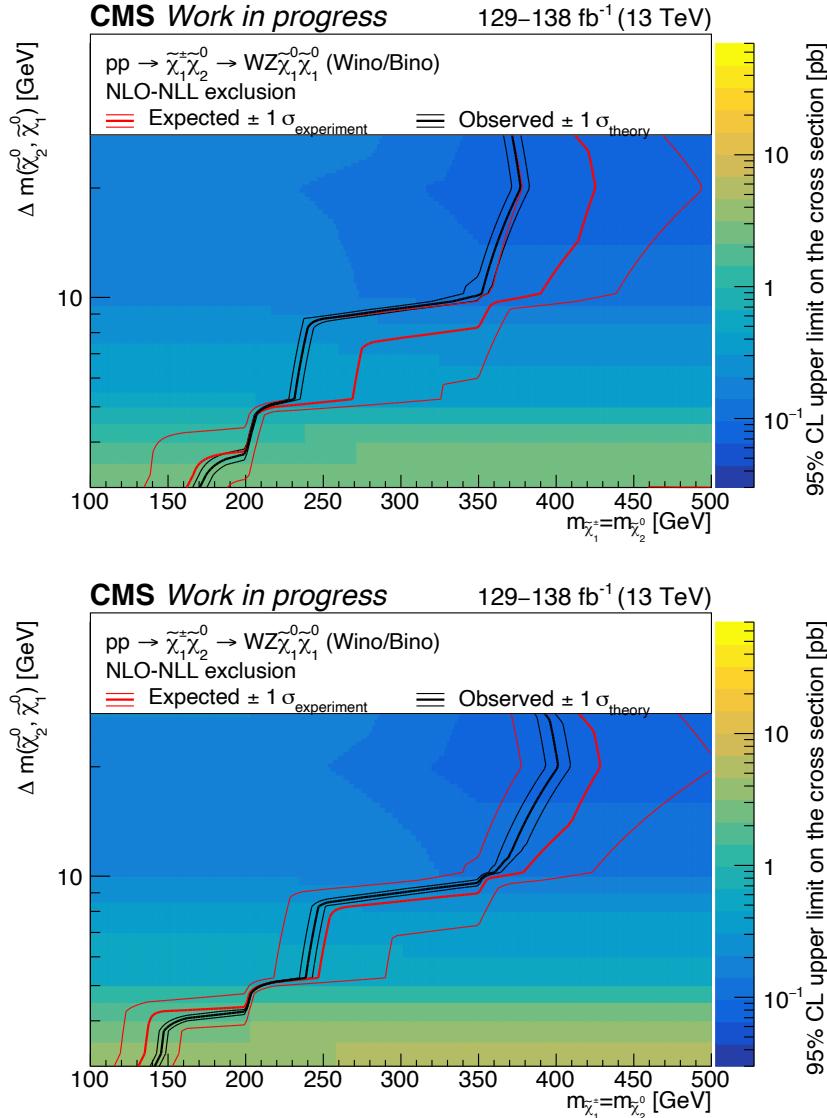


Figure 5.38: Observed and expected exclusion limits at 95% confidence level, based on the simplified model for chargino-neutralino production with $c\tau(\tilde{\chi}_2^0) = 10$ mm (top) and 100 mm (bottom), and assuming NLO+NLL wino production cross-sections. At the time of writing, minor components related to the fake lepton background and signal modeling uncertainties are still being implemented (see Section 5.8).

in the range of roughly 10 to 100 mm, due to relatively high signal efficiency and low SM background yields. The maximum observed exclusion is obtained for $c\tau = 100$ mm at a $\tilde{\chi}_2^0$ mass of 400 GeV and $\Delta m = 20$ GeV, which is an increase of approximately 120 GeV with respect to the scenario of promptly decaying $\tilde{\chi}_2^0$. For higher displacements, the sensitivity tends to decrease due to the lower reconstruction efficiency, as shown in Figure 5.21. Towards lower displacements, the sensitivity is affected by the presence of more SM background. However, for these signals (in particular $c\tau = 0.1$ mm) the combination of the prompt and displaced search regions is of large benefit, as signal events may end up in either of the search categories.

5.10 Conclusion

This search targets Supersymmetry realizations with compressed mass-spectra using final states with two or three soft leptons and p_T^{miss} . It is based on the full data-set recorded by the CMS experiment during Run 2, consisting of 138 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV. The search follows in the footsteps of Ref. [89], expanding the analysis reach in two directions. Firstly, the acceptance to more compressed signal scenarios is greatly increased by the usage a new, alternative *LowPT* electron reconstruction, which facilitates p_T thresholds as low as 1 GeV. Results are interpreted in the context of chargino-neutralino production in the wino-bino scenario. Limits are set on the sparticle masses, extending in the case of promptly decaying $\tilde{\chi}_2^0$ down to a mass-splitting of $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) = 900$ MeV at $m_{\tilde{\chi}_2^0} = m_{\tilde{\chi}_1^{\pm}} = 100$ GeV.

This result constitutes an important milestone in compressed electroweakino searches. For both the wino-bino and higgsino scenarios, experiments at the LEP collider were able to exclude chargino masses up to roughly 100 GeV [230] irrespective of the mass-splitting. At the LHC, searches from both CMS [89] and ATLAS [210] with soft lepton signatures have provided bounds that are significantly more stringent, except for mass-splittings below approximately 2 GeV. However, both CMS and ATLAS have recently performed searches with disappearing track signatures in the context of compressed higgsinos, resulting in the exclusion of mass-splittings below 0.4 GeV [97, 231]. Additionally, a recent search from ATLAS used low-momentum mildly-displaced tracks to exceed the LEP limits between mass-splittings of 0.4 and 0.9 GeV [232]. The search presented in this chapter closes the sensitivity gap between mass-splittings of 0.9 and 2 GeV, implying that – for the first time since LEP – searches at the LHC cover the full

range of electroweakino mass-splittings.

Besides targeting more compressed signal scenarios with the LowPT electron reconstruction, the analysis employs a new search category with soft displaced muon pairs to target – for the first time in CMS – mass-compressed *and* long-lived new physics realizations. Limits are set as function of the mass-splitting, $\tilde{\chi}_2^0$ mass and life-time at 95% confidence level. The maximum observed exclusion is obtained for $c\tau = 100$ mm at $\Delta m = 20$ GeV, where the $\tilde{\chi}_2^0$ mass exclusion of 400 GeV constitutes an increase of approximately 120 GeV compared to the case of promptly decaying $\tilde{\chi}_2^0$. To further extend the current reach in both prompt and displaced scenarios, the HL-LHC data-set is essential. This would greatly increase the statistical precision, which is currently still the limiting factor of searches for mass-compressed electroweakinos.

Chapter 6

Combination of Searches for Electroweak Supersymmetry

By the end of LHC Run 2 (2016-2018), the CMS experiment collected 137 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13 \text{ TeV}$ and set stringent limits on potential realizations of Supersymmetry at the electroweak scale [77–98]. The sensitivity of searches for new physics is typically limited by the amount of available data, but a significant increase of the total integrated luminosity is expected only during LHC Phase 2 (from 2029 onwards). Therefore, search efforts are instead focusing on maximally exploiting the available data by improving analysis techniques and performing statistical combinations of searches that target different production or decay channels within the same signal model.

This chapter presents the legacy Run 2 combination of searches for electroweak Supersymmetry with the CMS experiment. The author was one of the main analyzers and this work has been published in Ref. [214]. A plethora of search channels are used that target both (semi-)leptonic and hadronic final states, and the combined analysis is interpreted in terms of a variety of simplified models of R-parity conserving electroweak Supersymmetry. The targeted models and selected search channels are described in Sections 6.2 and 6.3, respectively. The combination strategy, including the orthogonality conditions and treatment of systematic uncertainty correlations are discussed in Section 6.4. Finally the combined results and interpretations are shown in Section 6.5 and compared against those of the component searches to highlight the additional signal parameter space excluded through combination.

6.1 Pushing the Limits with Statistical Combination

Since the start of Run 1 (2010-2012), experiments at the LHC have performed extensive search programs for new physics beyond the SM, with Supersymmetry amongst the top interests. By the end of Run 2 (2016-2018), the CMS experiment collected 137 fb^{-1} of proton-proton (pp) collision data at $\sqrt{s} = 13 \text{ TeV}$ and – in the absence of significant deviations between data and background expectations – set stringent limits on potential realizations of Supersymmetry at the electroweak scale [77–98] (see also Section 2.3). While Supersymmetry *might* reside at higher energy scales, development of collider technologies are needed to probe these scales and will not be available in the near future. It should be noted however, that Supersymmetry might still reside at the electroweak scale through realizations that are experimentally harder to probe. Search strategies are therefore frequently improved to increase the signal acceptance to more unconventional signatures, such as shown in Chapter 5. Furthermore, search sensitivities are often capped by statistical uncertainties, which will not be notably decreased until the High Luminosity LHC upgrade. As such, it is crucial to maximally exploit the available data via statistical combinations.

By design, individual searches are often optimal for only a limited region of the signal parameter space and target only a subset of the potential production and decay channels through which new physics could manifest itself. However, a combination of multiple search channels targeting the same signal scenario is expected be more sensitive by construction of the statistical methods that are used to interpret the analysis results. As shown in Appendix A, the statistical methods rely on the ratio of likelihoods corresponding to the signal+background and background-only hypotheses. When performing a combination, and under the condition that the search regions of the component analyses are strictly orthogonal, the likelihood functions of the individual searches are multiplied to produce a combined likelihood function. To illustrate the effect, consider two hypothetical searches that target different decay channels, where one has high sensitivity and one has low sensitivity (due to eg. different SM background populations). In setting upper limits on signal cross-sections, the likelihood of the latter goes to unity, which means that the combined limit simply tends to that of the more sensitive search. However, if both searches have sensitivity, the multiplication of the likelihood functions yields more stringent limits. This is the crucial benefit of statistical combinations.

This chapter presents the legacy Run 2 combination of searches for electroweak

Supersymmetry with the CMS experiment, based on all 137 fb^{-1} of pp collision data collected during Run 2. Each of the component searches (described in Section 6.3) targets unique final states and has reported upper limits on cross-sections of their respective signal hypotheses [86, 88–91, 233]. While the chosen input analyses and model interpretations constitute only a subset of the CMS search program, this selection facilitates a rich variety of potential combinations that could lead to increased signal sensitivity, due to the parameter space targeted by each search. Signal hypotheses with more exotic signatures (eg. long-lived decays) are less widely probed and hence not considered in this combination. Lastly, a combination of CMS analyses at $\sqrt{s} = 13 \text{ TeV}$ was performed before, but with only a partial data-set (35.9 fb^{-1} , taken in 2016) [234].

6.2 Signal Models

The combined search is interpreted in terms of four SMS models of Supersymmetry for the electroweak production of charginos, neutralinos and sleptons, as described in the following. In each model the lightest (typically two or three) Supersymmetric particles are assumed to be light enough to be produced at the LHC, while the other sparticles are much heavier and do not affect the observable phenomenology. All models correspond to R-parity conserving Supersymmetry realizations, implying that the sparticles are produced in pairs and each decay to a stable LSP – a potential thermal relic dark matter candidate – in association with a massive SM boson. The models and searches are characterized by the combination of SM particles emitted in these NLSP decays. For example, the production of $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$, which respectively decay to a W and Z boson (plus the LSP $\tilde{\chi}_1^0$), corresponds to the WZ topology.

The signal cross-sections for models of pair-produced electroweakinos are computed with minimal mixing of the gauge eigenstates, such that the mass eigenstates are bino-like, wino-like, or higgsino-like multiplets. Furthermore, the model of charged slepton production corresponds to a choice of mixing angles in the slepton mass-squared matrix that yields mass-degenerate $\tilde{e}_{L,R}$ and $\tilde{\mu}_{L,R}$.

The wino-bino model

This model describes the production of mass-degenerate wino-like NLSP $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$, as shown in Figure 6.1. The $\tilde{\chi}_1^\pm$ decays as $\tilde{\chi}_1^\pm \rightarrow W \tilde{\chi}_1^0$. The $\tilde{\chi}_2^0$ may decay to either a SM Z or H boson, in association with the LSP $\tilde{\chi}_1^0$, in order to allow for mixing of the neutral higgsino and gaugino states. As such, results are interpreted

Combination of Searches for Electroweak Supersymmetry

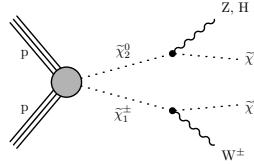


Figure 6.1: Production of $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$, considered in the wino-bino interpretation. The $\tilde{\chi}_1^\pm$ decays to a W boson and a $\tilde{\chi}_1^0$, and the $\tilde{\chi}_2^0$ decays to either a Z or H boson and a $\tilde{\chi}_1^0$.

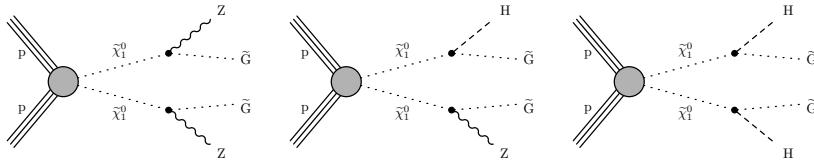


Figure 6.2: Production of $\tilde{\chi}_1^0\tilde{\chi}_1^0$ in the higgsino GMSB model. The $\tilde{\chi}_1^0$ particles each decay to a \tilde{G} and a SM gauge boson: (left) both $\tilde{\chi}_1^0$ decay to Z , (middle) one $\tilde{\chi}_1^0$ decays to Z and the other to H , and (right) both $\tilde{\chi}_1^0$ decay to H . Soft fermions from cascade decays of nearly mass-degenerate neutralinos and charginos are omitted.

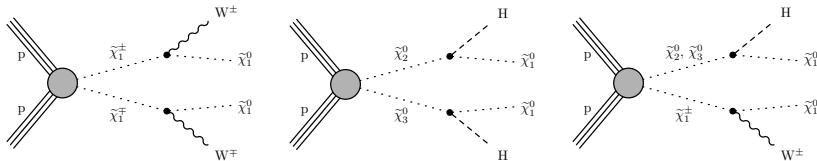


Figure 6.3: Production and decay modes considered for the higgsino-bino model, showing (left) the production of a pair of charginos that both decay to a W boson and the LSP, (middle) the production of a pair of neutralinos that both decay to H bosons and the LSP, and (right) the production of chargino-neutralino pairs followed by the decay of the chargino (neutralino) to a W (H) boson and the LSP.

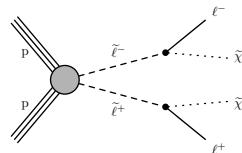


Figure 6.4: Direct slepton pair production, with each slepton decaying to a charged lepton and a $\tilde{\chi}_1^0$ LSP.

assuming branching ratio $\mathcal{B}(\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0) = 100\%$ or $\mathcal{B}(\tilde{\chi}_2^0 \rightarrow H\tilde{\chi}_1^0) = 100\%$ or with decays in equal measure.

The higgsino-bino and GMSB models

Two of the signal models are characterized by the production of light higgsinos. The higgsinos constitute a quadruplet of (nearly) mass-degenerate NLSPs – the lightest chargino pair $\tilde{\chi}_1^\pm$ and two neutralinos – while the LSP is either a gravitino or a bino-like $\tilde{\chi}_1^0$.

In the first higgsino scenario, motivated by a specific GMSB model [235], the LSP is the effectively massless gravitino \tilde{G} . The NLSPs are quasi-degenerate higgsinos, hence any combination of the lightest charginos and neutralinos can be produced¹: $\tilde{\chi}_1^\pm\tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm\tilde{\chi}_2^0$, $\tilde{\chi}_1^\pm\tilde{\chi}_1^\pm$ or $\tilde{\chi}_1^0\tilde{\chi}_2^0$. However, since the \tilde{G} coupling is suppressed by inverse orders of the breaking scale, the $\tilde{\chi}_1^0$ is metastable². Therefore, the aforementioned production modes all result in cascade decays to $\tilde{\chi}_1^0$ pairs, effectively yielding only $\tilde{\chi}_1^0\tilde{\chi}_1^0$ production, as shown in Figure 6.2. Furthermore, the higgsino mass-splittings are assumed to be very small, such that the SM products from the cascade decays are too soft to affect observables. As such, the event kinematics are fully determined by the decays of the $\tilde{\chi}_1^0$ pairs themselves, which may occur as $\tilde{\chi}_1^0 \rightarrow Z\tilde{G}$ or $\tilde{\chi}_1^0 \rightarrow H\tilde{G}$.

In the second higgsino scenario it is assumed that $|M_1| \simeq |\mu| \ll |M_2|$. Consequently, the $\tilde{\chi}_1^0$ is instead a bino-like LSP, and the NLSP higgsino multiplet consists of $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_3^0$. By choosing $|M_1| \simeq |\mu|$ one allows a (minimal) higgsino-bino mixture LSP, which can reduce the overabundance of thermally produced WIMPs that would occur for a pure-bino LSP, in order to match the observed dark matter relic density – yielding the *well-tempered neutralino* – without affecting the naturalness of model [236]. Furthermore, in contrast to the more conventional choice $|\mu| \ll |M_{1,2}|$, the mass-splitting $\Delta m(\text{NLSP,LSP})$ in this scenario is larger. As a result, the SM decay products are more energetic, making the production of this signal topology easier to detect. This model was not considered in the previous CMS combination [234]. The production modes are: $\tilde{\chi}_1^\pm\tilde{\chi}_2^0$, $\tilde{\chi}_1^\pm\tilde{\chi}_3^0$, $\tilde{\chi}_1^\pm\tilde{\chi}_1^\pm$ or $\tilde{\chi}_2^0\tilde{\chi}_3^0$ and it is assumed that $\mathcal{B}(\tilde{\chi}_1^\pm \rightarrow W\tilde{\chi}_1^0) = \mathcal{B}(\tilde{\chi}_2^0 \rightarrow H\tilde{\chi}_1^0) = 100\%$. The corresponding WW, HH and WH topologies are shown in Figure 6.3.

¹However, it can be shown that the process $pp \rightarrow \tilde{\chi}_i\tilde{\chi}_j$ is only non-vanishing for $i \neq j$ [204], hence the production of $\tilde{\chi}_1^0\tilde{\chi}_1^0$ and $\tilde{\chi}_2^0\tilde{\chi}_2^0$ is ignored.

²All searches nevertheless assume the promptly decaying $\tilde{\chi}_1^0$.

The slepton model

The last interpretation corresponds to direct production of charged slepton pairs with LSP $\tilde{\chi}_1^0$, as shown in Figure 6.4. The model assumes that $\tilde{e}_{L,R}$ and $\tilde{\mu}_{L,R}$ are mass-degenerate, while the staus are decoupled. In particular for this topology is the mass-compressed phase-space largely unprobed by the CMS Collaboration. However, in this combination this is addressed by a new interpretation of the analysis that targets soft multi-lepton final states, as will be described in Section 6.3.1.

6.3 Individual Searches

For the combined interpretation, a total of six input analysis are considered [86, 88–91, 233]. While each analysis targets final states with a notable amount of transverse missing energy p_T^{miss} (induced by the stable, undetected LSP), they impose unique selections on the visible hadronic and/or leptonic decay products, as described in the following.

6.3.1 (Semi)leptonic Final States

The “2/3 ℓ soft” analysis

This is the same search as described in Section 5.2, that targets mass-compressed chargino-neutralino production with the WZ decay topology in final states with two or three soft light leptons and p_T^{miss} [89]. The p_T of selected leptons ranges from 3.5 (5) GeV for muons (electrons) up to 30 GeV, and events must satisfy $p_T^{\text{miss}} > 125$ GeV, which in signal events is typically induced by the presence of initial state radiation. The main backgrounds consist of nonprompt leptons in eg. W+jets events, and of Drell-Yan and $t\bar{t}$ production. Control regions are employed to extract the normalization of the Drell-Yan, $t\bar{t}$ and WZ backgrounds from data.

The key variable in this analysis is the mass of the opposite-sign same-flavor (OSSF) lepton pair. In signal events this typically corresponds to the Z^* -boson mass, and hereby serves as a proxy for the mass-splitting between $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$. Conversely, the mass-splitting governs the $m_{\ell\ell}$ shape and bounds the spectrum from above at $m_{\ell\ell} = \Delta m$. Low mass SM resonances are explicitly vetoed, hence for the remaining background events the leptons typically originate from different mother particles. As such, the $m_{\ell\ell}$ variable is a great handle to discriminate signal from background. To optimally exploit this discrimination power, a parametric signal extraction is developed, consisting of unique $m_{\ell\ell}$ binnings of SR events for

each mass-splitting. The binnings were derived from analytical signal $m_{\ell\ell}$ shapes as described in Section 5.2.2.

Besides the wino-bino model, the “2/3 ℓ soft” analysis also provides a new interpretation for slepton pair production with small mass-splittings between the sleptons and the LSP $\tilde{\chi}_1^0$. The compressed spectrum of smuon or selectron pairs has not been studied before by the CMS Collaboration. In this topology the sleptons decay to a pair of soft opposite-sign same-flavor leptons and p_T^{miss} , similar to the WZ topology. As such, the analysis has decent acceptance to compressed sleptons without modifying the search strategy. This new interpretation is therefore based on the same background estimations, analysis regions and systematic uncertainties. Only the 3 ℓ search regions are excluded, since it does not contribute to the signal acceptance.

The final likelihood fit for the slepton model relies on a modified use of the parametric signal extraction. Instead of $m_{\ell\ell}$, the SRs are binned in $m_{\text{T}2}(k, k, m_\chi)$, defined as:

$$m_{\text{T}2}(k, k, m_\chi) = \min_{\vec{p}_T^{\text{miss}(1)} + \vec{p}_T^{\text{miss}(2)} = \vec{p}_T^{\text{miss}}} \left[\max \left(M_{\text{T}}^{(1)}(m_\chi), M_{\text{T}}^{(2)}(m_\chi) \right) \right] \quad (6.1)$$

This variable is a measure for the mass M of particles, produced in pairs, that each decay to semi-visible final states [237, 238]. The minimization introduces two dummy momenta, in this case representing the unknown momenta of the true neutralinos χ , which together make up the total observed \vec{p}_T^{miss} in the event. The $M_{\text{T}}^{(i)}$ are the transverse masses constructed from $\vec{p}_T^{\text{miss}(i)}$ and either of the visible objects k , in this case the leptons. For simplicity, the neutralino mass m_χ is fixed to 100 GeV, which is expected to have a negligible effect on the search sensitivity compared to a dynamic value of m_χ . The distribution of the $m_{\text{T}2}(\ell, \ell, 100)$ variable ranges from m_χ to M , so from 100 to 130 GeV, given that the leptons must satisfy $p_T < 30$ GeV. A set of parametric SR binnings are designed separately for each p_T^{miss} bin of the 2 ℓ SR and each $(m_{\tilde{\ell}}, m_{\tilde{\chi}_1^0})$ mass-hypothesis. As for the WZ topology, the design strategy enforces equal signal quantiles in each bin, while ensuring the background estimation in each bin satisfies $N_{eq} > 1.5$. However, in the absence of an analytical description of the signal $m_{\text{T}2}$ shape, the expected yields from simulated MC samples are used instead. An example of the resulting post-fit SR distributions $m_{\text{T}2}(\ell, \ell, 100)$ for the mass-hypothesis $(m_{\tilde{\ell}}, m_{\tilde{\chi}_1^0}) = (125, 115)$ is shown in Figure 6.5.

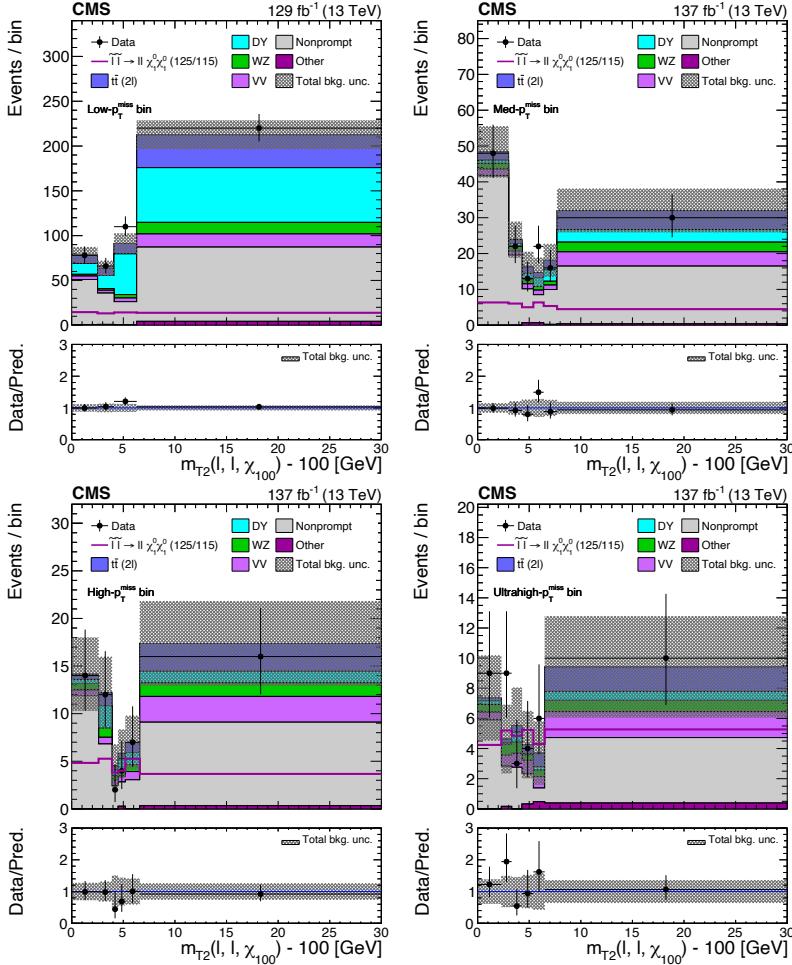


Figure 6.5: Post-fit distributions of the $m_{T2}(\ell, \ell, m_\chi)$ variable, with m_χ fixed to 100 GeV [214]. The figures correspond to the low- (upper left), medium- (upper right), high- (lower left), and ultrahigh- (lower right) p_T^{miss} bins of the “2 ℓ soft” signal region of Ref. [89]. The distributions are based on the parametric binnings derived for the mass-point $m_{\tilde{\ell}} = 125$ GeV, $m_{\tilde{\chi}^0} = 115$ GeV, and the corresponding pre-fit signal distribution is overlaid for illustration (purple line). Note that the signal distribution is approximately flat across m_{T2} by construction of the parametric binning procedure.

The “ 2ℓ on-Z” and “ 2ℓ nonresonant” analysis

From the search of Ref. [233] two search categories are included in this combination, referred to as “ 2ℓ on-Z” and “ 2ℓ nonresonant”. The first category targets the wino-bino and GMSB models, while the second targets the slepton model, both in the uncompressed region of the model parameter space. For both categories, events are required to have exactly two OSSF leptons (electrons or muons) and p_T^{miss} . The lower bounds of leading (trailing) lepton p_T are fixed at 25 (20) GeV, and selections on the p_T^{miss} , $m_{\text{T2}}(\ell, \ell, \tilde{\chi}_1^0)$ and $m_{\text{T2}}(\ell b, \ell b, \tilde{\chi}_1^0)$ are chosen to reduce SM background. The remaining backgrounds consists primarily of Drell-Yan, $t\bar{t}$ and associated production of a Z boson and neutrinos (eg. $t\bar{t}Z$, $WZ \rightarrow 3\ell\nu$, $ZZ \rightarrow 2\ell 2\nu$), which are estimated from data in CRs.

The “ 2ℓ on-Z” category targets on-shell production of Z bosons by requiring $86 < m_{\ell\ell} < 96$ GeV. This category is further divided into three SRs, each requiring an additional hadronically decaying higgs or vector (V) boson; In the resolved VZ SR events must have at least two jets and no b-tagged jets. In the Lorentz-boosted VZ SR events must have < 2 AK4 jets and ≥ 1 AK8 jet. In the HZ SR there must be 2 b-tagged jets consistent with the Higgs boson mass.

The “ 2ℓ nonresonant” category targets slepton pair production and requires $20 < m_{\ell\ell} < 65$ or $m_{\ell\ell} > 120$ GeV. This category is further divided into two SRs, based on the presence of jets in addition to the lepton pair.

The “ $2\ell\text{SS}/\geq 3\ell$ ” analysis

This analysis selects a variety of final states with at least two leptons and p_T^{miss} [86]. It targets the WZ, WH, HH, HZ and ZZ signal topologies and hence is used here for the interpretations with the wino-bino, higgsino-bino and GMSB models. Twelve orthogonal search categories are defined:

- Category SS: a pair of same-sign light leptons and no τ_h .
- Category A: three light leptons, including an OSSF lepton pair.
- Category B: three light leptons, with no OSSF lepton pairs.
- Categories C, D, E, and F: three leptons, at least one of which is a τ_h
- Categories G, H, I, J, and K: more than three leptons.

The categories are further divided via kinematic cuts on eg. particle p_T , $m_{\ell\ell}$, and the scalar p_T sum H_T of selected jets; details can be found in Ref. [86]. The SS

category is the only region with exactly two leptons, and the main backgrounds consist of nonprompt leptons or charge mismeasurements, which are estimated from data in CRs. Categories A, B, G and H require three or four light leptons (ie. no τ_h), which allows the estimation of the main backgrounds (WZ and ZZ production) directly from MC samples, validated in low- p_T^{miss} CRs. The remaining categories, C-F and I-K, require at least one τ_h and are dominated by Drell-Yan and $t\bar{t}$ production with associated nonprompt τ_h candidates, estimated from data using the tight-to-loose method.

Category A has the highest acceptance to signal with the WZ decay topology, but also suffers from the highest backgrounds among all categories. Therefore, while the other categories rely on a cut-based approach with standard kinematic variables, category A relies on an alternative approach for the WZ interpretation, that is based on parametric neural networks in order to construct variables with higher discriminating power. For the other interpretations a cut-based version of category A is used.

The “1ℓ2b” analysis

This search [88] targets the WH topology and is used here for the wino-bino and higgsino-bino interpretations. The final state consists of two b-tagged jets compatible with Higgs decays ($90 < m_{bb} < 150$ GeV), one light lepton with $p_T > 30$ GeV and p_T^{miss} . Events with Lorentz-boosted Higgs bosons are selected by requiring a single large-radius jet ($R = 0.8$), identified by the DEEPAK8 algorithm [239]. SM backgrounds with a single leptonically decaying W boson are suppressed by requiring $m_T(\ell, p_T^{\text{miss}}) > 150$ GeV, and the remaining backgrounds are predominantly from processes with two leptonically decaying W bosons (eg. $t\bar{t}$ and tW production). These SR contributions are estimated by applying transfer factors on data yields in CRs. SR events are further categorized based on the presence of a boosted H-tagged jet, p_T^{miss} , and the number of AK4 jets.

6.3.2 All-hadronic Final States

The “4b” analysis

This analysis [90] is used here for the GMSB and higgsino-bino interpretations. It targets $\text{HH} \rightarrow 4\text{b}$ decays, distinguishing between resolved and boosted signatures, corresponding to whether the b quarks from each H boson decay are contained in two separate AK4 jets or a single AK8 jet, respectively. Charged leptons are

vetoed and at least $p_T^{\text{miss}} > 150$ GeV is required.

In the resolved case, SR events must have four or five AK4 jets with $p_T > 30$ GeV and are further categorized depending on the number b-tagged jets. The two H candidates are formed by building pairs out of the four jets with highest b tag score and minimizing the difference of m_{bb} . The difference must be less than 40 GeV and the average compatible with Higgs decays ($100 < \langle m_{bb} \rangle < 140$ GeV).

In the boosted case, at least two AK8 jets with $p_T > 300$ GeV are required and the two highest p_T jets must be loosely compatible with the Higgs mass ($60 < \langle m_J \rangle < 260$ GeV). Events are then categorized based on the number of double b-tagged jets.

The main background consists of $t\bar{t}$ production, as well as W or Z production in association with jets, all of which are estimated from data in dedicated CRs.

The “Hadr. WX” analysis

The second analysis selecting all-hadronic final states [91] targets decays of the WW, WZ and WH signal topologies and is used for the wino-bino and higgsino-bino interpretations. The final states include at least two AK8 jets, $p_T^{\text{miss}} > 200$ GeV, $H_T > 300$ GeV and no charged leptons. The dominant backgrounds are from W+jets, Z+jets, and $t\bar{t}$ production, which are estimated with data in control regions. The analysis makes extensive use of the tagging capabilities of the DEEPAK8 algorithm [239] to distinguish AK8 jets from different SM boson decay modes, including the mass-decorrelation of the neural network tagging scores. The mass-decorrelated version of the DEEPAK8 W tagger (referred to as the V tagger) is used to tag hadronic decays of both W and Z bosons, while the W tagger without mass-decorrelation (referred to as the W tagger) is used to tag only W bosons, but with lower misidentification rate.

A b-veto SR is defined by the absence of b-tagged jets and is sensitive to both WW and WZ signal topologies. At least one AK8 jet must be tagged by the W tagger, at least one other by the V tagger, and the AK8 jet mass must satisfy $65 < m_J < 105$ GeV.

A b-tag SR is defined by the presence of at least one b-tagged jet and is sensitive to both WZ and WH topologies. AK8 jets from Z and H bosons are tagged using the DEEPAK8 bb tagger with mass decorrelation, and the jet mass must satisfy $75 < m_J < 140$ GeV. This SR is further categorized into three regions: One with at least one jet tagged as W candidate and bb-tagged H candidate (WH); One with at least one tagged W boson candidate and a veto on jets tagged as Higgs candidates (W); And one with at least one Higgs boson candidate and a veto on

jets tagged as W candidate (H).

6.4 Combination Strategy

An overview of the signal models and topologies targeted by each of the input analysis is shown in Table 6.1. For each of the signal models a simultaneous maximum likelihood fit is performed with all the relevant signal and control regions of the input analyses. To compute the combined likelihood by multiplications it is strictly required that the analyses regions be orthogonal. Therefore, when a combination of any two searches results in (partially) overlapping analysis regions, orthogonality is recovered either by excluding the overlapping regions from the fit or modifying the region definitions as described in Section 6.4.1. The fit itself is performed with the statistical framework COMBINE, described in Ref. [240]. It builds the likelihood as a product of Poisson probability density functions in all regions of the observed data yields given the expected background and signal yields. The signal contribution is scaled by a positive valued strength parameter μ that is free in the fit. Systematic uncertainties are implemented as nuisance parameters in the fit, and common systematic uncertainties are correlated among the input analyses as described in Section 6.4.2.

Search	Wino-bino		GMSB			Higgsino-bino			Sleptons $\ell^+\ell^-$
	WZ	WH	ZZ	ZH	HH	WW	HH	WH	
2/3 ℓ soft [89]	all								2 ℓ soft
2 ℓ on-Z [233]		EW		EW	EW				
2 ℓ nonres. [233]									Slepton
2 ℓ SS/ $\geq 3\ell$ [86]	SS, A(NN)	SS, A-F	all	all	all			SS, A-F	
1 ℓ 2b [88]		all						all	
4b [90]					all		3-b, 4-b, 2-bb		
Hadr. WX [91]	all	all				ex H		ex H	

Table 6.1: Summary of the searches considered in the combination and the analysis categories that contribute to the interpretation of each signal model and topology. The following notations are used: For the “2 ℓ on-Z” analysis, “EW” refers to the resolved and boosted VZ SRs and the HZ SR. For the “2 ℓ nonresonant (nonres.)” search, “Slepton” refers to the two dedicated slepton SRs, those requiring $N_{\text{jet}} = 0$ and $N_{\text{jet}} > 0$. For the “2 ℓ SS/ $\geq 3\ell$ ” search, “A(NN)” indicates SR A with the parametric neural network signal extraction. For the “Hadr. WX”, “ex H” denotes all SRs except the b-tag H SR.

6.4.1 Orthogonality of Search Regions

For the possible combinations summarized in Table 6.1, the object and event selections of all signal and control regions included in the likelihood fit were scrutinized to identify potential overlaps. Also, the fraction of observed data events shared between any two analyses was determined by comparing the events populating these signal and control regions. Out of all possible cases of overlap, most were found to be below the sub-percent level, and thus considered acceptable given their negligible effect on the final results. In two cases, however, the overlap was notable, and required a customised strategy to be resolved.

Firstly, for the higgsino-bino interpretation, the combination of the “Hadr. WX” and “4b” analyses results in overlap for multiple sets of signal and control regions, due to the selection of events with at least two AK8 jets and subsequent b-tag requirements. Orthogonality is recovered by removing low purity regions from the combination, being the b-tag H SRs and CRs of the “Hadr. WX” analysis and the single bb-tag SRs and CRs of the boosted “4b” category, which are the least sensitive for their targeted signal topologies. While for the individual searches this reduces the expected 95% confidence level upper limit on the signal cross-section (tested with various mass-hypotheses) by the order of a few percent, it maximizes the combined sensitivity for the WH and HH signal topologies.

Secondly, for the wino-bino interpretation, overlap was found between regions of the “ $2/3\ell$ soft” and “ $2\ell SS/\geq 3\ell$ ” analyses, corresponding to the selection of events with 3 light leptons. In particular, the 3ℓ WZ CR of the “ $2/3\ell$ soft” search, used to constrain the WZ background normalization with data, almost fully overlaps with the category A of the “ $2\ell SS/\geq 3\ell$ ” search, the main SR for the WZ interpretation with sensitivity also to the WH topology. Therefore, the overlap was resolved by removing the 3ℓ WZ CR from the combined fit. To instead constrain the WZ background normalization in the “ $2/3\ell$ soft” analysis with the observed yields in the “ $2\ell SS/\geq 3\ell$ ” search, a common nuisance parameter is assigned that correlates the normalization of these backgrounds between both analyses.

Besides the 3ℓ WZ CR, also the 3ℓ SR of the “ $2/3\ell$ soft” search overlaps, with both categories A and B of the “ $2\ell SS/\geq 3\ell$ ” search, due to lepton p_T selections that are not fully orthogonal. The upper bound of the lepton p_T is 30 GeV in the “ $2/3\ell$ soft” analysis, whereas the requirements on the minimum p_T of electrons and muons vary between 10 and 25 GeV in the aforementioned categories of the “ $2\ell SS/\geq 3\ell$ ” analysis. However, none of these regions could be removed without a price of sensitivity. Given that these regions are important

for the mass-compressed signal scenarios, for which the “ $2/3\ell$ soft” analysis has the highest sensitivity, the overlap was removed by raising the leading lepton p_T threshold in the “ $2\ell SS/\geq 3\ell$ ” search to 30 GeV. For consistency, this modification applies to all model interpretations, not just those for which the “ $2/3\ell$ soft” and “ $2\ell SS/\geq 3\ell$ ” searches are combined. The consequences of this modification were carefully estimated. In category A, for the WZ CR of the “ $2\ell SS/\geq 3\ell$ ” roughly 1% of background events were lost. For the CR used to evaluate the modeling of asymmetric photon conversions 10% of the yields were lost, while the agreement between data and expected background is maintained at a similar level. In most SRs, the expected background is reduced by less than 1%. Only in regions with $m_{\ell\ell} < 75$ GeV, the background yields are reduced by up to 7%. Similarly, yields of uncompressed signal hypotheses are largely unaffected, while the compressed signal models suffer a loss of 5-7% in yields. In category B, signal and background contributions are reduced by roughly 5% and 1-2%, respectively. In terms of the expected upper limits, the sensitivity for a wide range of mass-hypotheses (tested for the GMSB and wino-bino models) was found to be mostly unchanged. Only in the compressed case of the WZ topology the expected sensitivity decreased by up to 10% for mass-splittings between 20 and 70 GeV. However, this loss of roughly 10 GeV in NLSP mass exclusion is compensated by the performance of the “ $2/3\ell$ soft” analysis.

6.4.2 Systematic Uncertainties

Treatments of the systematic uncertainties for each of the input analyses are described in the respective publications [86, 88–91, 233]. These uncertainties are implemented as nuisance parameters in the likelihood fit. The scheme for correlating uncertainties in the combination of analyses is summarized in Table 6.2.

In general, uncertainties with sources of the same nature are treated as correlated. This includes uncertainties on the measurement of the total integrated luminosity, modeling of the pileup distributions and correction factors accounting for a gradual shift in the timing of trigger information from the ECAL, present during the 2016 and 2017 data-taking periods. Also uncertainties related to the modeling of object efficiencies and initial state radiation (ISR) are fully correlated. Only for the slepton and wino-bino interpretations, the lepton efficiencies are treated as uncorrelated between the “ $2/3\ell$ soft” and “ 2ℓ nonresonant” analyses as they cover disjoint regions of the model parameter space.

The uncertainty in the modeling of the trigger efficiency is partially correlated,

6.4 Combination Strategy

specifically between analyses that share primary high-level trigger paths. Also the uncertainties related to the renormalization and factorization scales μ_R and μ_F are partially correlated, since the associated nuisance parameters of some input analyses only affect the total yields, while others affect only the shape of simulated MC distributions or both.

Uncorrelated are all statistical uncertainties accounting for the size of the simulated MC samples. Also the dominant uncertainties on the background normalizations are uncorrelated, since these tend to be either statistical in nature or are related to the estimation methods themselves. Finally, uncertainties beyond those in Table 6.2 are analysis-specific and treated as uncorrelated as well.

Source	Correlated?
General	
MC sample size	No
SM background normalization	No
Integrated luminosity	Yes
Trigger efficiency	Partially
Pileup	Yes
Trigger timing	Yes
Objects and signal modeling	
Lepton efficiency	Yes
Jet energy resolution	Yes
Jet energy scale	Yes
b (mis)tagging efficiency	Yes
AK8 bb tagging efficiency	Yes
AK8 jet mass resolution	Yes
μ_R and μ_F	Partially
ISR modeling	Yes
Attributable to the CMS fast simulation	
p_T^{miss} modeling	Yes
b (mis)tagging	Yes
AK8 bb tagging	Yes
AK8 bb mass	Yes

Table 6.2: Sources of systematic uncertainties and the correlations between analyses. For the SM background normalizations, all uncertainties are uncorrelated, except for the WZ normalization, which is correlated between the “ $2\ell SS/\geq 3\ell$ ” and “ $2/3\ell$ soft” searches. Furthermore, lepton efficiencies are generally correlated, except between the “ $2/3\ell$ soft” and “ 2ℓ nonresonant” searches, because they cover disjoint regions of the model parameter space.

6.5 Results and Interpretations

The combined maximum likelihood fit is performed for each mass-hypothesis and signal model. No significant deviations from the SM background are observed, consistent with the results from the individual searches. Cross section upper limits at 95% CL as a function of the sparticle masses are set using a modified frequentist approach, employing the CL_s criterion and an asymptotic formulation [240–243] (see also Appendix A). The upper limits on the model cross sections and the mass exclusion contours corresponding to $\mu = 1$ are summarized in the following.

For the wino-bino model, the combined analysis results are shown in Figure 6.6 and compared with the individual analysis results in Figure 6.7. Three scenarios for the $\tilde{\chi}_2^0 \rightarrow H\tilde{\chi}_1^0/Z\tilde{\chi}_1^0$ branching ratios are considered. For the WZ topology a substantial expansion of the excluded parameter space is obtained by the combined analysis, in particular for the uncompressed regions. NLSP masses below 875 GeV are excluded for a light LSP, and LSP masses below 420 for a 700 GeV NLSP. These are increases of roughly 125 GeV with respect to the individual searches. Sensitivity for the compressed region (top-right plots of Figures 6.6 and 6.7) is driven by the “ $2/3\ell$ soft” analysis, with a cross-over to the “ $2\ell SS/\geq 3\ell$ ” analysis around $\Delta m = 30$ GeV. In this cross-over regions, the observed exclusion limits are weaker than expected due to small data excesses in both analyses, as already reported in the respective publications [86, 89]; These occur at $20 < m_{\ell\ell} < 30$ GeV in the ultra-MET bin of the 2ℓ SR and at $10 < m_{\ell\ell} < 30$ GeV in the low-MET bin of the 3ℓ SR of the “ $2/3\ell$ soft” analysis, and at regions with low $m_{T2}(\ell, \ell)$ or high $m_{T2}(\ell, \ell)$ and low $p_T(\ell, \ell)$ in category SS of the “ $2\ell SS/\geq 3\ell$ ” analysis. The observed exclusion amounts to roughly 2 standard deviations with respect to the expected exclusion. For the WH topology the excluded parameter space is expanded as well, but less than expected due to small ($\simeq 1$ standard deviation) data excesses in several p_T^{miss} bins of the WH SR in the “Hadr. WX” analysis. This is visible in the mixed WZ-WH topology as well. Finally the analyses with the best exclusion limit for each mass-hypothesis are shown in Figure 6.8.

For the higgsino GMSB model, the results as function of the $\tilde{\chi}_1^0$ mass are shown in Figure 6.9, corresponding to decays with $\mathcal{B}(\tilde{\chi}_1^0 \rightarrow H\tilde{G}) = 100\%$, $\mathcal{B}(\tilde{\chi}_1^0 \rightarrow Z\tilde{G}) = 100\%$ or $\mathcal{B}(\tilde{\chi}_1^0 \rightarrow H\tilde{G}) = \mathcal{B}(\tilde{\chi}_1^0 \rightarrow Z\tilde{G}) = 50\%$. The results for the full range of possible branching ratios are shown in Figure 6.10. Mass values for the $\tilde{\chi}_1^0$ are excluded below 840, 1025 and 760 GeV for the ZZ, HH and 50% mixed topology, respectively. The combined result expands the excluded region by up to 100 GeV at $\mathcal{B}(\tilde{\chi}_1^0 \rightarrow H\tilde{G}) = 40\%$.

The results for the higgsino-bino interpretation are shown in Figure 6.11. This model was not considered in the previous combined search for electroweak Supersymmetry by CMS [234]. The mass-values for degenerate $\tilde{\chi}_2^0$, $\tilde{\chi}_3^0$ and $\tilde{\chi}_1^\pm$ are excluded between 225 and 800 GeV for $\tilde{\chi}_1^0$ masses below 50 GeV.

Lastly, the limits for slepton pair production are shown in Figure 6.12. The excluded parameter space is not expanded by a combined analysis, since the “2/3 ℓ soft” and “2 ℓ nonresonant” searches cover disjoint regions of the model parameter space. However, for the first time in CMS, the compressed region (right plot of Figure 6.12) is targeted. Here the “2/3 ℓ soft” analysis excludes slepton masses up to 215 GeV at $\Delta m = 5$ GeV.

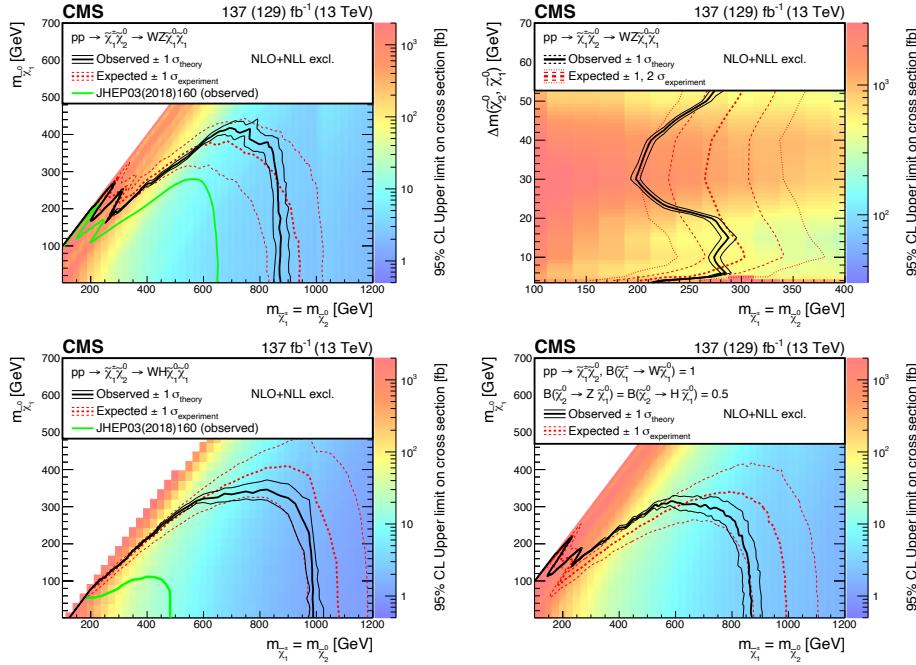


Figure 6.6: Cross section limits and exclusion boundaries for chargino-neutralino production in the wino-bino model [214]. The figures correspond to the WZ topology for the full parameter space (upper left) as well as the compressed region (upper right), the WH topology (lower left), and the mixed topology with 50% branching fraction to WZ and WH (lower right). For some signal regions the analysis was based on a subset of the data, corresponding to a total integrated luminosity of 129 fb^{-1} .

6.6 Conclusion

Six previously reported searches for Supersymmetry are combined to provide the legacy Run 2 results on possible manifestations of electroweak Supersymmetry, using 137 fb^{-1} of proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$ recorded with the CMS detector. No significant deviation from the SM expectation has been observed, and limits are set on the production of charginos, neutralinos and sleptons based on simplified SUSY models.

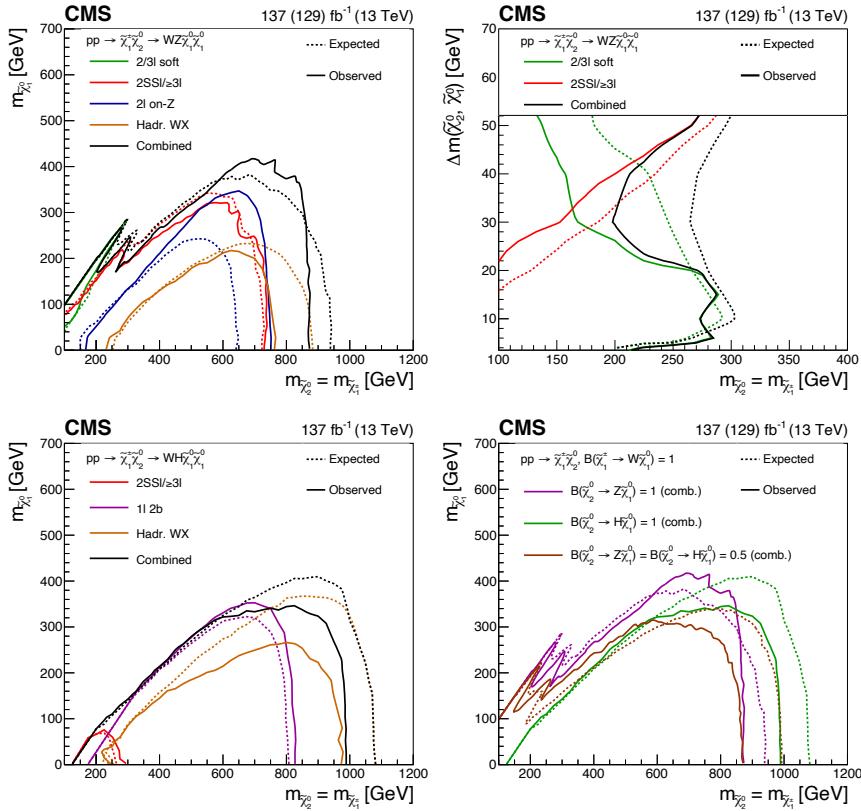


Figure 6.7: Exclusion boundaries for chargino-neutralino production in the wino-bino model from the individual analyses [214]. The figures correspond to the WZ topology for the full parameter space (upper left), the corresponding compressed region (upper right), and the WH topology (lower left). The combined contours (as in Figure 6.6) for these topologies are also shown. In the lower right figure the combined contours for these and the mixed topology are overlaid. For some signal regions the analysis was based on a subset of the data, corresponding to a total integrated luminosity of 129 fb^{-1} .

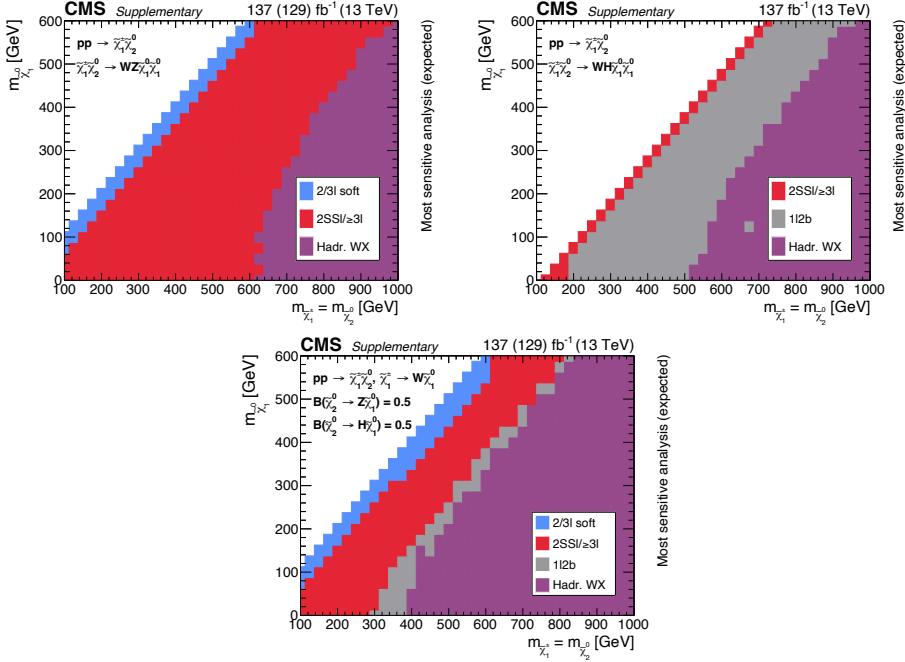


Figure 6.8: Analysis with the best expected exclusion limit at each point in the $(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ mass-plane for chargino-neutralino production in the wino-bino model, corresponding to the WZ topology (upper left), WH topology (upper right) and the mixed topology with 50% branching fraction to WZ and WH (bottom) [214]. For some signal regions the analysis was based on a subset of the data, corresponding to a total integrated luminosity of 129 fb^{-1} .

The combined analysis yields a considerable expansion of the excluded parameter space, compared to those excluded by any of the individual searches, as well as the previous CMS combination [234]. In general, wino-like chargino masses are excluded up to 990 GeV, and higgsino-like neutralino masses in the GMSB model are excluded up to 1025 GeV. The mass-bounds are increased by 100–510 GeV with respect to the Ref [234], and the excluded parameter space is expanded by as much as 125 GeV with respect to the most sensitive component searches.

The higgsino-bino model was not considered in Ref. [234]. With the combined legacy Run 2 results, mass-generate higgsino-like $\tilde{\chi}_2^0$, $\tilde{\chi}_3^0$ and $\tilde{\chi}_1^\pm$ are excluded between 225 and 800 GeV for a bino-like $\tilde{\chi}_1^0$ with $m_{\tilde{\chi}_1^0} < 50$ GeV.

Lastly, the “2/3 ℓ soft” component analysis was reoptimized and – for the first time in CMS – performed a dedicated search for slepton pair production targeting the mass-compressed model parameter space. In the combined search, slepton

Combination of Searches for Electroweak Supersymmetry

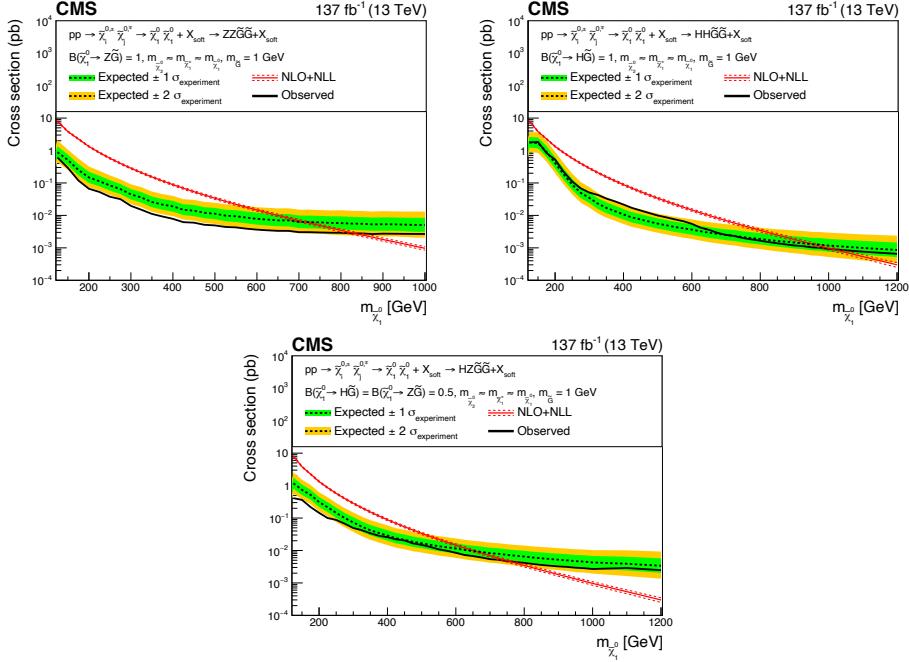


Figure 6.9: Expected and observed exclusion limits for the higgsino GMSB model, considering neutralino-neutralino production with the ZZ topology (upper left), the HH topology (upper right), and the mixed topology with 50% branching fraction to H and Z (lower) [214]. Masses below the crossing point of the theory cross-section and observed upper limit are excluded.

masses up to 215 GeV are excluded at $\Delta m(\tilde{\ell}, \tilde{\chi}_1^0) = 5$ GeV, and in the uncom-pressed region between 130 and 700 GeV for $m_{\tilde{\chi}_1^0} < 50$ GeV.

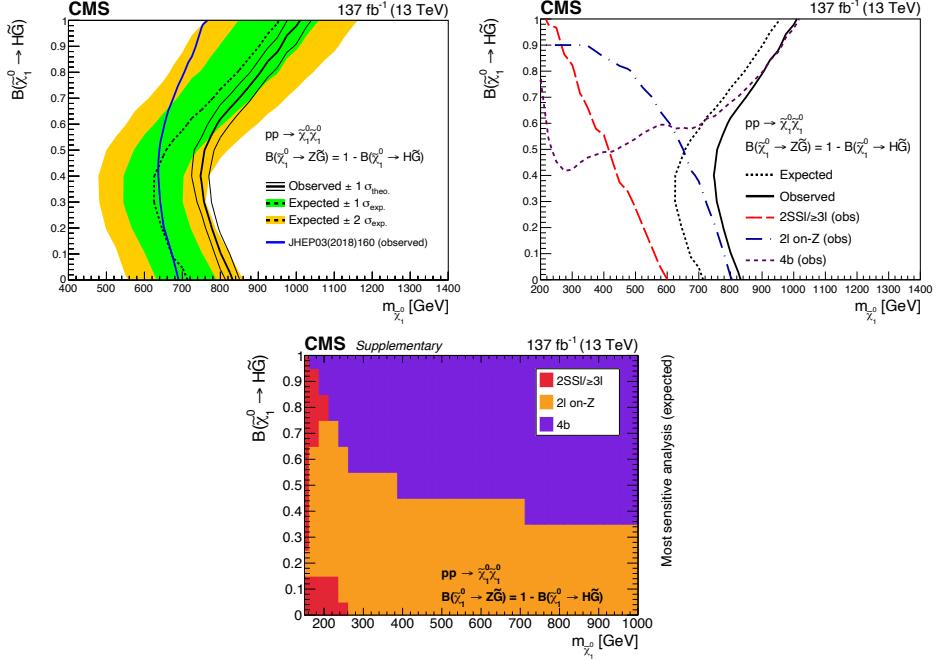


Figure 6.10: NLSP mass exclusion limits for neutralino-neutralino production in the higgsino GMSB model as a function of the $\mathcal{B}(\tilde{\chi}_1^0 \rightarrow H\tilde{G})$ branching fraction [214]. Upper left: expected and observed limits for the combination of the searches, compared with the observed limits of the combination [234] based on the 2016 CMS data. Upper right: expected and observed exclusion limits for the combination in comparison with those of the component searches. Bottom: analysis with the best expected exclusion limit for each signal hypothesis.

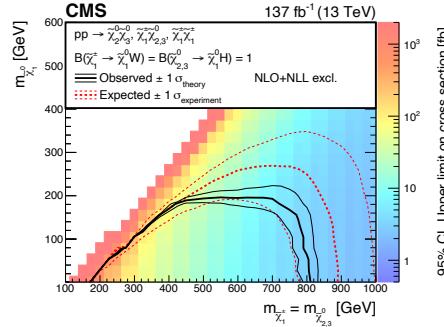


Figure 6.11: Cross section limits and the expected and observed mass exclusion boundaries for the higgsino-bino model [214].

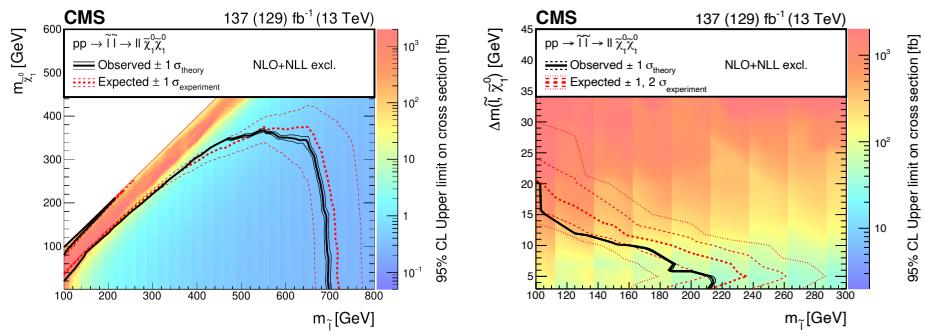


Figure 6.12: Cross section limits and the expected and observed mass exclusion boundaries for the the full mass plane from the combination (left), and for the compressed region from the “2/3 ℓ soft” search (right) [214]. For some signal regions the analysis was based on a subset of the data, corresponding to the integrated luminosity of 129 fb^{-1} .

Chapter 7

Level-1 Trigger Algorithms for Electrons at the HL-LHC

As described in Chapter 3, the LHC will be upgraded to the High Luminosity LHC (HL-LHC) during 2026-2028. The HL-LHC is designed to deliver 14 TeV proton-proton (pp) collisions at a factor of 5 to 7.5 times the nominal LHC luminosity, ie. up to $7.5 \times 10^{34} \text{ cm}^2 \text{s}^{-1}$, which is crucial for many statistically limited measurements of Standard Model properties and searches for new physics, such as those shown in Chapters 5 and 6. However the running conditions of the HL-LHC pose serious challenges for the experiments.

The amount of simultaneous pp interactions per bunch crossing will increase by a factor 5 with respect to the average at the end of LHC Run 2, resulting in many more particles traversing the detector volumes. One of the main challenges will be the online selection of collision events. The CMS Collaboration has completely redesigned the CMS Level-1 Trigger in order to cope with these harsh conditions, while maintaining the physics acceptance of the Run 2 system. To further improve the trigger performance, novel algorithms are being developed, that maximally exploit the more granular information from sub-detector upgrades and new trigger capabilities.

Sections 7.2 and 7.1 summarize the main components of the Phase 2 upgrade of the CMS Level-1 Trigger and the baseline electron trigger algorithms, respectively, as presented the Technical Design Report of 2020 [148]. The work of the author is shown in Section 7.3 onwards, consisting of the development of an alternative electron reconstruction and identification strategy for the upgraded Level-1 Trigger, that brings notable improvements to the performance with respect to the baseline algorithm. This work has been published Ref. [244].

7.1 The Phase 2 Upgrade of the CMS Level-1 Trigger

By the end of LHC Run 3 in 2026, the LHC will have delivered roughly 400 fb^{-1} of collision data to the ATLAS and CMS detectors. This will mark the end of Phase 1 and the beginning of a new era: Phase 2. The LHC will be upgraded to the High Luminosity LHC (HL-LHC) and start colliding protons at a center-of-mass energy $\sqrt{s} = 14 \text{ TeV}$ with roughly 7 times higher instantaneous luminosity compared to that of the LHC [143]. Over the course of its lifetime, the HL-LHC is expected to deliver a grand total of 4000 fb^{-1} of collision data to the experiments, a factor 10 increase with respect to Phase 1. This will drastically improve the statistical precision for exploring the electroweak scale. However, the running conditions of the HL-LHC pose serious challenges for the detectors. The average amount of simultaneous pp interactions per bunch crossing (pileup, PU) will increase to roughly 200 (compared to 40 by the end of Run 2). This causes many more particles to traverse the detector volumes, increasing the risk of radiation damage, misidentification of particles and mismeasurements of particle momenta and energies. Therefore, a plethora of detector upgrades is required to operate the CMS experiment in the HL-LHC environment, as is described in Chapter 3.

One of the main challenges will be the online selection of collision events, performed by the hardware based Level-1 Trigger (L1T). The Phase 1 L1T makes trigger decisions based on partial event information from the Calorimeter and Muon systems, reducing the event rate from 40 MHz to 100 kHz. Studies based on simulation have shown that using todays L1T algorithms in the HL-LHC conditions, would result in a trigger rate of 4000 kHz, far above the bandwidth allowed by the data acquisition system [148]. In the past several years, the CMS Collaboration has therefore completely redesigned the L1T system, in order to maintain and extend the physics acceptance of the Run 2 system also in the HL-LHC era, while keeping the rates within a manageable bandwidth. The preliminary design of the CMS Phase 2 L1T upgrade has been documented in a Technical Design Report in 2020 [148].

A schematic of the design architecture is shown in Figure 7.1. One of the most anticipated upgrades to the L1T is the *Track Trigger* (depicted in green). For the first time in CMS, charged particle tracks within the Outer Tracker volume ($|\eta| < 2.4$) will be reconstructed at the pp collision rate of 40 MHz by the electronic boards of the Track Finder (TF) as part of the detector backend. The so-called “trigger primitives” (described in more detail in Section 7.2.3) will be sent to the

7.1 The Phase 2 Upgrade of the CMS Level-1 Trigger

Global Track Trigger (GTT), which computes tracker-only based objects, such as jets and missing transverse momentum. Additionally, the GTT will reconstruct the position of collision vertices, and will therefore serve as an important handle for PU mitigation.

Also the *Calorimeter Trigger* (depicted in red) will benefit notably from improved trigger primitives generated in the backend electronics. In particular, in the end-cap region ($|\eta| > 1.52$) the new High Granularity Calorimeter (HGCAL) will provide a rich and detailed description of particle showers, which can be used for multivariate particle identification. In the barrel region ($|\eta| < 1.52$), the energy measurement granularity will increase by a factor 25 due to the improved readout of the existing Electromagnetic Calorimeter.

A central role is reserved for the *Correlator Trigger* (depicted in yellow), another new addition to the Phase 2 L1T design. Since the physical implementation of the trigger architecture makes extensive use of state-of-the-art Field Programmable Gate Arrays (FPGA) and processors, connected via high speed optical links, it will be possible to process data from multiple subdetectors on the same electronics board. The Correlator Trigger will receive inputs from the Calorimeter, Muon and Track Trigger systems and can execute more sophisticated algorithms

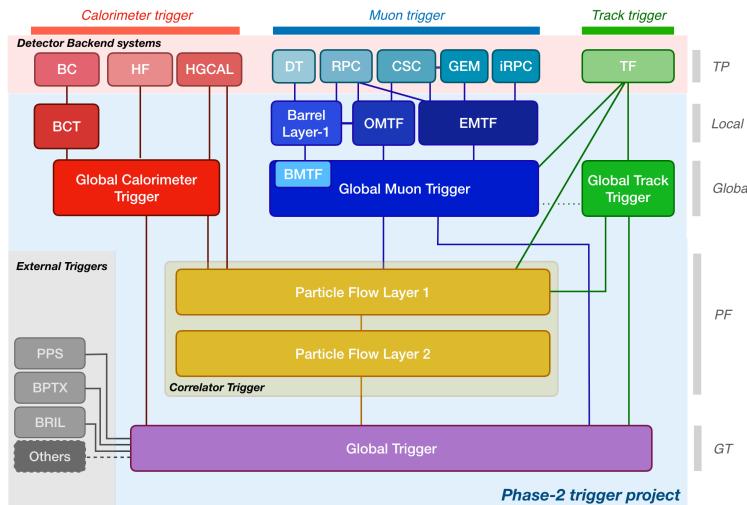


Figure 7.1: The CMS L1T architecture for Phase 2 [148]. The highly modular design provides optimum flexibility and robustness, and can be divided into four data processing paths: The Calorimeter trigger (red), the Muon trigger (blue), the Track trigger (green) and the Correlator trigger (yellow).

to produce higher-level trigger objects. A key feature of the Correlator Trigger will be the ability to geometrically match Tracker objects to the “stand-alone” Calorimeter deposits and Muon objects. This can greatly reduce combinatorial background and improve the momentum measurements of the matched objects, and therewith decreases the trigger rate. A further rate reduction may be achieved via track-based pile-up resilient isolation requirements. To illustrate the sizable rate reduction achieved via track-matching and track-based isolation, an example for electron candidates is shown in Figure 7.2 (left).

Besides track-matching stand-alone Calorimeter and Muon objects, another valuable feature of the Correlator Trigger is the capability to host global event reconstruction techniques, such as Particle-Flow (PF) [149] and Pileup Per Particle Identification (PUPPI) [203]. The Particle-Flow reconstruction algorithm, described in Chapter 4, is already widely used in CMS for offline analysis and in the High Level Trigger. However, since it strongly relies on efficient charged particle track reconstruction and high granularity information from the calorimeters to resolve close-by particles, this was not implemented in the Phase 1 L1T. Furthermore, the algorithm was designed to run on CPUs, processing candidates sequentially and with complex inter-dependencies, making it particularly time consuming in high occupancy events. To run Particle-Flow in the Phase 2 L1T, us-

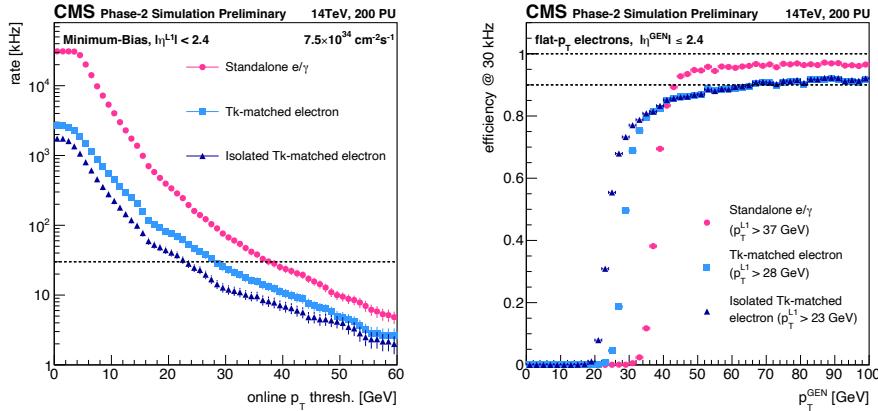


Figure 7.2: Physics performance of various electron objects used in the CMS Phase 2 L1T menu [244]. Standalone e/γ (pink) corresponds to clusters identified with calorimeter-only information. These may additionally be track-matched (light blue) and pass track-based isolation requirements (dark blue). Left: Single electron trigger rates as function of the Level-1 reconstructed (cluster) p_T threshold. Right: Identification efficiency as function of generator-level electron p_T corresponding to a fixed rate of 30 kHz, which is the typical rate budget for single electron triggers.

ing the new TF inputs and highly granular information from the Calorimeters, a simplified version of the original algorithm was designed. While PF is known for its holistic event interpretation that combines information from multiple sub-systems, the association of reconstructed tracks and energy deposits is inherently *local* (mostly performed with ΔR comparisons). The simplified PF version for the L1T benefits from this fact by dividing the detector volume into multiple (partially overlapping) rectangular areas in the η - ϕ plane, referred to as PF regions. This facilitates parallel reconstruction of individual particles using only the detector inputs in the vicinity of the candidate. With reduced complexity and parallelization, it is possible to fit the algorithm within the resource and latency constraints of the L1T.

Major changes to the upgraded L1T and data acquisition hardware also facilitate an increase of L1T output rate from 100 kHz to 750 kHz, which will partially cover the expected increase of trigger rate in Phase 2. Furthermore, the maximum latency, i.e. the maximum time allotted to make the trigger decision, will be increased from $3.8\ \mu\text{s}$ to $12.5\ \mu\text{s}$. This allows the usage of more complex trigger algorithms that need more time to be evaluated during data-taking, such as HG-CAL clustering, track finding, and Particle-Flow. Combined with the increased processing power of FPGAs - the building blocks of the L1T hardware - and high inter-connectivity of the L1T architecture, this also provides a stage for more advanced machine learning based trigger strategies.

Machine learning techniques in high energy physics are typically employed for e.g. particle identification and jet clustering, applying sophisticated multivariate selections to exploit correlations between objects and features. However, due to their generally high computational demands, machine learning techniques have seen only limited use in the online event selection. The challenge of the L1T task in Phase 2, together with the recent gains in FPGA processing power, has inspired the development of new compiler packages based on High Level Synthesis (HLS), referred to as `HLS4ML` and `CONIFER`, that allow neural network (NN) based algorithms and boosted decision trees (BDT) respectively, to be ported to FPGAs [245, 246]. Thanks to these new tools it is therefore possible to construct low latency trigger algorithms, that can fit within the available resources of modern FPGAs. Triggers based on machine learning techniques will maximize discrimination power for signal objects and event topologies in the overwhelming background, hereby increasing the overall physics acceptance and reducing the trigger rate compared to the traditional cut-based trigger algorithms.

7.2 Track-matched Electron Reconstruction in the CMS Endcaps

This section describes the electron reconstruction from the generation of trigger primitives to the identification of the final electron object that is used to make trigger decisions. The identification of the final electron object, the “*TkElectron*”, has been considered as the baseline identification algorithm ever since its introduction in the Phase 2 L1T Technical Design Report in 2020 [148]. It relies on the new track-matching capabilities of the Correlator Trigger and high granularity shower shape information provided by HGCAL in order to achieve sufficient rate reduction to maintain the Phase 1 p_T -thresholds of single electron triggers also in Phase 2. The baseline algorithm is described below. From Section 7.3 onwards, a new alternative identification strategy is introduced to further improve the trigger performance for electrons in the CMS end-caps. Therefore, the description of the baseline algorithm in the following will focus on electrons in the forward regions ($|\eta| > 1.5$) of the detector.

7.2.1 HGCAL Trigger Primitive Generation

The electron reconstruction starts with building the trigger primitives from energy deposits in the HGCAL, performed by the calorimeter trigger primitive generator (HGCAL-TPG), see Figure 7.3. At the pp collision rate of 40 MHz, raw input data is grouped into trigger cells by on-detector electronics, summing the energy deposits of neighbouring sensors. In the silicon section of HGCAL these trigger cells have a granularity of roughly 4 cm^2 . In the scintillator section the trigger cells are groups of scintillator tiles spanning roughly 4 to 10 cm in both the azimuthal and radial direction. Since not all data from each trigger cell can be sent out, a dynamic energy threshold is defined based on the trigger cell occupancy and only the selected trigger cells are sent to the off-detector backend electronics. The energy deposits in unselected trigger cells are summed in groups of roughly 48 trigger cells (module sums) and sent downstream as well.

After repacking and calibrating the data (not explicitly shown in Figure 7.3), the backend electronics will perform a energy clustering procedure. The selected trigger cells are used to fill histograms projecting the energy deposits in the $r/z - \phi$ plane. After reducing fluctuations with a 1D smoothing kernel, local energy maxima are identified. By attaching close-by trigger cells to these so-called “seeds”, 3D clusters are built that are the main trigger primitives sent from the HGCAL to

7.2 Track-matched Electron Reconstruction

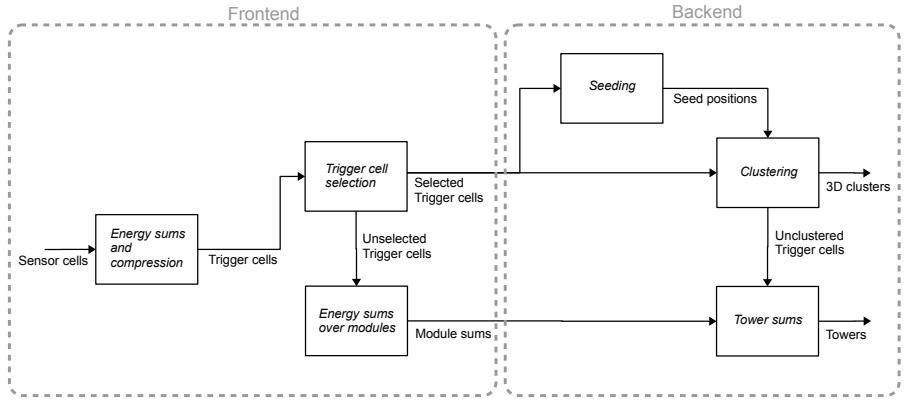


Figure 7.3: Schematic diagram of the processing steps of the HGCAL trigger primitive generation [148]. The final trigger primitives are a set of 3D clusters and trigger towers that are sent to the L1T boards.

the central L1T. The choice of histogram binning, smoothing, seed definition and cluster size all contribute to a delicate compromise between shower containment and contamination from low energy pileup deposits.

Finally, the unclustered trigger cells (backend) are combined with the energy sums from unselected trigger cells (frontend) into projective towers that have a

Data type	Content	Size / BX (bits)
Header	Total energy, BX, $N_{clusters}$	32×6
Cluster	Position	32
	Minimal energies	32
	Minimal shape quantities	40
	Flags (quality and data format)	24
	Optional energies	32
	Optional shape quantities	128
	Optional substructure quantities	128
Total clusters ($\times N_{clusters}$)	Minimal	128
	Maximal	416
Tower	Energy	12
	EM fraction	4
Total 2600 towers		43 k

Table 7.1: Data format of trigger primitives generated by the HGCAL-TPG [148].

size of $\pi/36$ in the $\eta - \phi$ plane. These towers contain all energy deposits not included in the 3D clusters, which could be useful e.g. for L1T jet reconstruction algorithms downstream. The set of 3D clusters and projective towers are sent in two copies, one to the Global Calorimeter Trigger and one to the Correlator Trigger. The data format of each copy is shown in Table 7.1. The bulk of the data string is reserved for detailed information of the reconstructed 3D clusters. Besides the cluster position, both an electromagnetic and hadronic interpretation of the cluster energy is included. Additionally, while the exact choice of feature definitions is still being finalized, the data string also contains a rich description of the cluster development in the longitudinal and transverse direction.

7.2.2 Stand-alone e/γ Identification

After generating the trigger primitives, two sequential identification procedures take place in order to determine the type of the particle that initiated the particle shower in HGCAL. Since in both these procedures only information from HGCAL is used, namely the 3D clusters, this step is referred to as “stand-alone” identification. Furthermore, both procedures make use of boosted decision trees (BDTs) to maximally exploit the discrimination power of the 3D cluster descriptions.

The first object identification procedure aims at distinguishing genuine e/γ and hadronic clusters, originating from hard interaction processes, from generally low energetic pile-up clusters. This algorithm is referred to as the “PU veto”. It serves as a first line of defense against the overwhelming background from PU, while maintaining adequate efficiency for genuine low p_T clusters. The BDT model is trained on a large set of HGCAL clusters, that are reconstructed in a mixture of simulated MC events with photons and pions, and are required to pass $5 < p_T < 60$ GeV and $1.55 < |\eta| < 2.85$. Clusters are labelled as signal if they are matched to a generator level photon or charged pion, satisfying $\Delta R < 0.1$ and $p_T^{\text{reco}}/p_T^{\text{gen}} > 0.3$. To account for the curvature of the charged pion trajectory in the magnetic field of the detector, the geometrical matching is performed with the generator objects propagated to the calorimeter surface. Background clusters are taken from a sample of simulated MC events with photons as well, but may not be matched to generator level objects, in order to obtain a sample of clusters from PU only. The BDT is trained with five longitudinal and lateral shower shape features, as shown in Appendix C. After training, a threshold on the BDT output score corresponding to 94% signal efficiency and 99% background rejection is chosen as working point (WP) for the selection. All 3D clusters constructed by

the HGCAL-TPG are required to pass this working point before being processed by the Correlator boards.

The second identification procedure is the “stand-alone e/γ ID”, consisting of another multivariate model which specifically serves to identify clusters initiated by electrons or photons and is implemented with BDTs as well. For the BDT training, signal clusters are defined as those passing $p_T > 10$ GeV and matched to a generator electron that satisfies $p_T > 20$ GeV and $\Delta R < 0.2$. Background clusters are defined as all clusters with $p_T > 20$ GeV and not matched to a generator electron. The input features are a set of 5 longitudinal and 4 lateral shower shape variables as shown in Appendix C. Separate models are trained for $1.5 < |\eta| < 2.7$ and $2.7 < |\eta| < 3.0$ as this was found to improve background rejection compared to a model trained in the inclusive η range. However, for the remainder of this chapter only the model trained for the range $1.5 < |\eta| < 2.7$ is relevant, given that the Level-1 tracks are reconstructed only up to $|\eta| < 2.4$. For this model, two working points are defined: a loose WP corresponding to 99.5% signal efficiency at 89% background rejection, and a tight WP with 97.5% signal efficiency at 94.2% background rejection.

HGCAL clusters that pass both the PU veto and the stand-alone e/γ ID (tight WP) are promoted to calorimeter-only e/γ candidates. These candidates are then used for a bremsstrahlung recovery step. As electrons propagate through the detector, they may interact with nuclei from the detector material and radiate a photon. This process (bremsstrahlung) can therefore result in one or more low energy clusters from photons in the vicinity of a high energy cluster from the electron. To recover the energy loss from bremsstrahlung, low energy clusters are therefore merged into the highest energetic cluster in a $\Delta\eta \times \Delta\phi = 0.02 \times 0.1$ window, improving the p_T response in particular at high pseudo-rapidity, where the amount of traversed detector material is larger. The final rate of stand-alone e/γ candidates as function of p_T threshold, as well as the reconstruction efficiency are shown in Figure 7.2.

7.2.3 Track Trigger Primitive Generation

The capability to use Tracker information in the L1T is one of the key upgrades that allows the trigger system to maintain the Phase 1 physics acceptance, even in the running conditions of the HL-LHC. Via new handles, such as vertex reconstruction, track isolation, and track-matching stand-alone objects from the Calorimeter and Muon systems, the Tracker information will play a major role in reducing

background rates. At the basis of these new capabilities stand the increased L1T latency and total bandwidth, as well as the increased processing power of modern FPGAs, which can be exploited in a (time) multiplexed way. Moreover, the CMS Tracker will be upgraded with the L1 Track Trigger requirements amongst the main drivers for the design of the Outer Tracker for Phase 2 [121].

In order to limit the sheer data volume resulting from track reconstruction with an average of 200 pile-up events, the Outer Tracker will make use of p_T *modules*, implemented in front-end on-detector electronics. As charged particles propagate through the magnetic field of the detector, their trajectories bend in the transverse plane, with the curvature depending on the particle p_T . The p_T modules are able to reject particles below a certain p_T threshold, by correlating the hits in

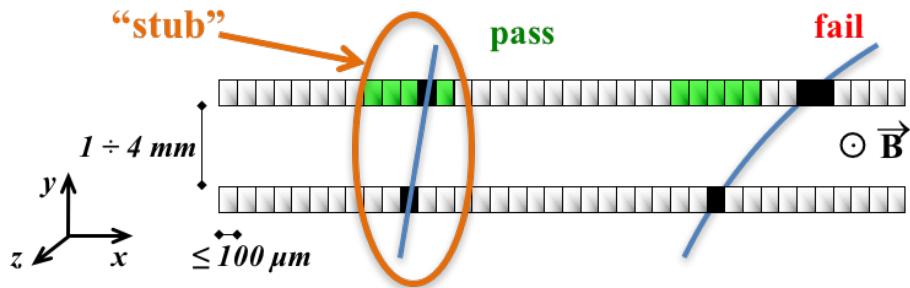


Figure 7.4: Schematic of a p_T module that creates stubs from pairs of Tracker hits [121]. The green boxes indicate the range of possible hits in the top layer that would result in a stub with $p_T > 2 \text{ GeV}$.

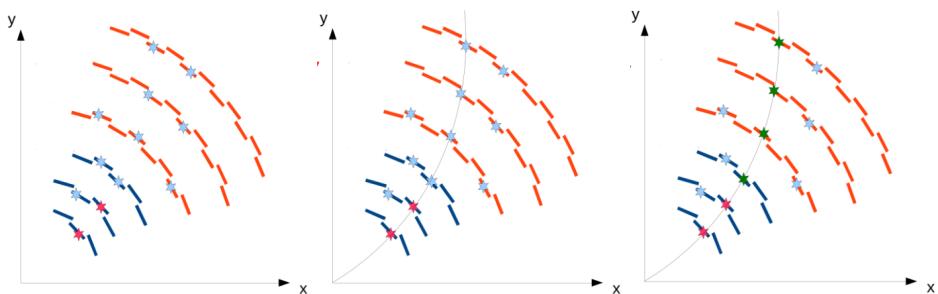


Figure 7.5: Schematic of the tracklet finding algorithm [247]. Left: two stubs, marked as red stars, are selected to seed the algorithm. Middle: Based on the selected stubs a track is projected to the other Tracker layers to search for matching stubs. Right: A final track fit is performed, including all stubs consistent with the track projection.

closeby sensors. As illustrated in Figure 7.4, a p_T module consists of two closely-spaced sensors within the same Tracker layer that are read out by a common set of application-specific integrated circuits (ASIC). These ASICs will combine hits from the two sensors layers to form *stubs*. The bending angle of these stubs is limited, resulting in an effective p_T threshold of 2 GeV. This in turn reduces the data volume sent to the backend electronics by a factor 10.

The off-detector Track Finder (TF) system receives the stubs at 40 MHz and constructs so-called *tracklets*, which are the trigger primitives used by the central L1T [247]. The formation of tracklets is illustrated in Figure 7.5. The first step is to form pairs of stubs that - together with the beamspot at $(x, y) = (0, 0)$ - are consistent with a track having $p_T > 2$ GeV, $|\eta| < 2.4$ and $|z_0| < 15$ cm. Based on these initial tracklet parameters, the tracklet is projected to the other Tracker layers in order to find consistent hits. All matching stubs are then used for a track fit to reconstruct the final track parameters. To speed up this track-finding algorithm, the Outer Tracker is divided in nine ϕ sectors, which are being processed in parallel by separate FPGAs. Furthermore, for each sector 18 FPGAs are reserved for time-multiplexing, in order to increase the maximum processing time from 25 ns to 450 ns.

The final data format sent from the TF to the central L1T is shown in Table 7.2. It contains all track fit parameters with enough precision to match the maximum resolution of the L1 Tracker. Furthermore, there are several bits reserved for the evaluation of multi-variate classifiers, which are still being developed and are to be implemented in the Track Finder boards.

Table 7.2: Data format of trigger primitives generated by the Track Finder [148].

Track parameter	Number of bits
q / R	15
ϕ	12
$\tan(\lambda)$	16
z_0	12
d_0	13
χ^2 / dof	4
bend- χ^2	3
hit mask	7
track quality MVAs	3
other quality MVAs	6
track isValid	1
spare	4
total	96

7.2.4 Track-matched Electron Identification

Section 7.2.2 described the stand-alone e/γ ID algorithm used to identify electrons from the HGCAL trigger primitives (the 3D clusters) only. As can be seen from Figure 7.2 (right), this algorithm yields a plateau efficiency close to 1. However, to keep the rate within the budget of roughly 30 kHz allotted for single electron triggers in the Phase 2 menu is $O(10)$ GeV higher¹ than the thresholds used in Phase 1. In other words, the stand-alone e/γ ID does not allow the Phase 1 physics acceptance to be maintained in the Phase 2 environment, which is the primary goal of the Phase 2 L1T upgrade.

To further reduce the rate, a final identification step is performed, that involves geometrically matching the HGCAL clusters – after having passed the standalone ID – to a L1 track. To account for the effect of the magnetic field on charged particle tracks, the tracks are propagated to the calorimeter surface and matched to the cluster by applying an elliptical cut in the $\eta - \phi$ plane, as is shown in Figure 7.6. The ellipse is extended in the ϕ direction to cover the track extrapolation uncertainty arising from p_T mismeasurements, while it is more narrow in the η direction to minimize background combinatorics. The tracks that are sent to the Correlator Trigger boards, where the matching procedure takes place, are required

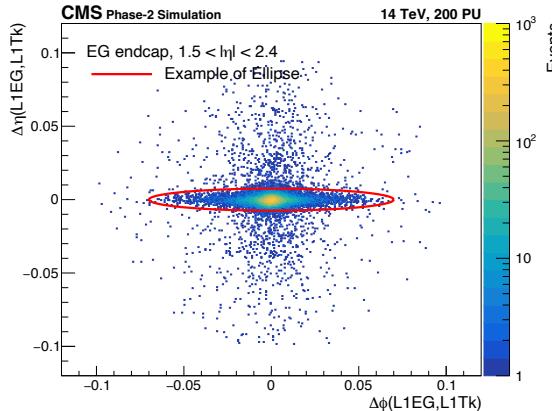


Figure 7.6: The spatial separation in the $\eta - \phi$ plane between HGCAL clusters and the closest L1 track [148]. The red line indicates the elliptical cut of the form $(\Delta\eta/0.0075)^2 + (\Delta\phi/0.07)^2 < 1$, which is used to construct track-matched electrons in the baseline identification algorithm.

¹The comparison is based on the p_T threshold of 37 GeV shown in Figure 7.2 and the L1 electron seeds – without isolation requirements – of the lowest-threshold unprescaled single electron HLT path used in 2018, that require at minimum $p_T > 26$ GeV.

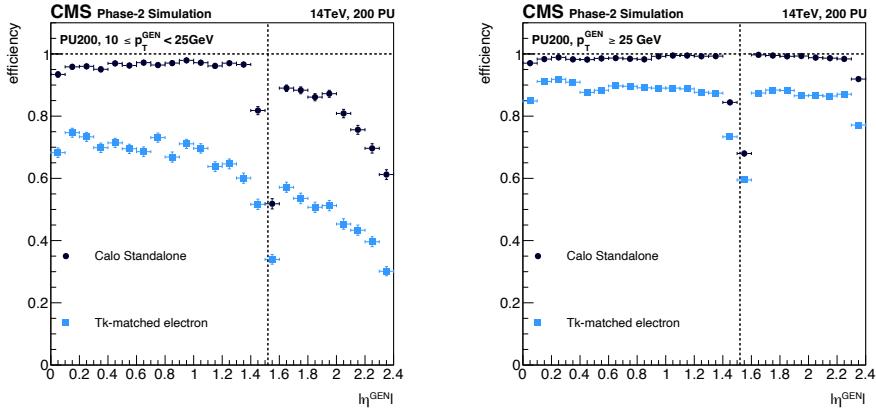


Figure 7.7: Efficiencies of identifying a L1 e/γ object within $\Delta R < 0.2$ of a generator level electron, as function of pseudo-rapidity. Electrons are generated with $10 < p_T^{\text{GEN}} < 25$ GeV (left) and $p_T^{\text{GEN}} > 25$ GeV (right) with an average of 200 PU events. The baseline TkElectron identification (light blue) shows decreased efficiency compared to the stand-alone e/γ ID (dark blue), in particular for low p_T objects in the CMS end-caps ($|\eta| > 1.5$).

to pass $p_T > 5$ GeV or ($p_T > 2$ GeV and $\chi^2 < 15$). That is, an additional cut on the Kalman Filter (KF) track-fit quality is applied to the softest tracks in order to reduce the object multiplicity and, consequently, the required matching resources that scale as the product of the number of tracks and clusters. Conversely, this χ^2 cut is relaxed for tracks with $p_T > 5$ GeV, in order prevent penalizing electron tracks, which tend to have worse fit quality, since their Bethe-Heitler like energy loss (dominated by bremsstrahlung) is poorly approximated by KF tracking².

The background rate from low energy PU candidates is mainly controlled by increasing the p_T cut on the tracks considered for the matching. To allow for similar trigger thresholds as the single electron triggers of Phase 1, the track p_T requirement is increased to > 10 GeV. The total rate reduction achieved by matching the standalone clusters to a track with $p_T > 10$ GeV is shown in Figure 7.2 (left).

This final L1 object, consisting of a track with $p_T > 10$ GeV that is matched to an HGCAL cluster that passed the PU and standalone e/γ ID, is referred to as the “TkElectron”. This is the baseline and *only* algorithm described in the Technical Design Report of the Phase 2 L1T [148], that allows the Phase 1 thresholds of single electron triggers to be maintained also at an average of 200

²As described in detail in Section 4.3.2, the offline reconstruction of electron tracks relies on fitting the Tracker hits with a Gaussian Sum Filter (GSF), which better models the energy loss from photon radiation than KF tracking. However, GSF tracking in the L1T is not possible due to resource constraints. The L1T therefore relies on KF tracking.

PU events. Even though this is clearly a major achievement, requiring a track with $p_T > 10$ GeV comes at a heavy price of decreased efficiency, which is shown in Figure 7.7. The track p_T requirement naturally affects softer objects the most: for electrons with p_T between 10 and 25 GeV, the efficiency in the end-caps is decreased by roughly 30%, while for electrons with p_T below 10 GeV the efficiency is effectively non-existent. Besides affecting the efficiency of low p_T electrons, the track requirement also affects the high p_T objects, decreasing the efficiency for electrons in the end-caps with $p_T > 25$ GeV by roughly 15%.

7.3 Improving Performance for Track-matched Electrons

Given the vital importance of electron triggers across the broad CMS physics program, the limited efficiency of the TkElectron strongly requests new identification algorithms that improve the electron trigger performance. The work presented in this section constitutes a new algorithm for electron reconstruction and identification that provides a major increase in selection efficiency for the same background rates, across the full p_T spectrum. It is referred to as the “Composite electron identification” and has been published in Ref. [244].

7.3.1 Composite Electron Identification

Before introducing the Composite electron identification strategy, it is important to stress that any alternative to the TkElectron has to be based on the same trigger primitive reconstructions and pre-selections described in Section 7.2. While the electron reconstruction in general could benefit from improvements in the TPs (cluster and tracks) targeted towards the complex case of e/γ objects, at the time of writing the baseline TP generation has not yet been fully demonstrated in firmware. As such, it is too early to consider optimizations for e/γ objects. Therefore, the baseline generation of trigger primitives in the Tracker and HGCAL is taken here to be fixed by design, as well as the object pre-selections that are imposed to satisfy hardware constraints of the detector backends and links to the central L1T.

The Composite electron identification (Composite ID) is a new algorithm, proposed to replace the baseline TkElectron identification of the Phase 2 L1T Technical Design Report [148]. It aims at identifying electrons using composite objects constructed from HGCAL clusters and L1 tracks in a more holistic way and fundamentally differs from the baseline identification in three aspects, as summarized in Table 7.3 and described in more detail below.

7.3 Improving Performance for Track-matched Electrons

First and foremost is the lower bound on the track p_T . The baseline identification has no handle to control the track quality, and instead relies solely on the track $p_T > 10$ GeV cut to reduce the rate from low energy PU and combinatorial background candidates. Consequently, the efficiency to identify genuine electrons from the hard interaction process is decreased as well. The only way to recover this efficiency loss is by relaxing the track p_T requirement. Therefore, the Composite ID imposes no explicit track p_T requirements beyond those already applied in the pre-selection of the track TP inputs.

The second difference is the geometrical matching of tracks and clusters. Figure 7.8 shows the angular distance between tracks and clusters in single electron events, simulated with an average of 0 PU events. For clusters with $p_T > 10$ GeV most often the electron track can be found within the elliptical matching window of the cluster, as used in the baseline identification. However, for clusters with $p_T < 10$ GeV, the angular distance at the calorimeter surface is generally too large. These cases occur when the electron - as it propagates through the Tracker volume - radiates a photon (bremsstrahlung). The photon will not be affected by the magnetic field of the detector and continues in a straight line to the calorimeter, initiating a HGCAL cluster. The kink in the electron track from photon radiation is not accommodated by the current KF-based tracking, hence the electron track might be mis-reconstructed. The problem gets worse in a high-PU environment, where tracker hits from PU particles might inadvertently be used in the reconstruction of the electron track. As a result the track-parameters for radiating electrons tend to be poorly estimated, which widens the angular separation between the track and cluster. It may therefore happen that the cluster will

Table 7.3: Overview of (pre-)selection and matching steps performed in the baseline TkElectron and Composite electron identification algorithms.

	TkElectron ID	Composite ID
Cluster preselection	$1.52 < \eta < 2.4$ $p_T > 1$ GeV Pass PU ID Pass loose WP of e/γ BDT	
Track preselection	$ \eta < 2.4$ $p_T > 10$ GeV	$p_T > 5$ GeV or ($p_T > 2$ GeV and $\chi^2 < 15$)
Geometrical track-cluster matching	$(\Delta\eta/0.0075)^2 + (\Delta\phi/0.07)^2 < 1$	$\Delta R < 0.2$
Algorithmic selection	Cluster passes tight WP of e/γ BDT	Composite candidate passes WP of Composite BDT

either be matched to *another* track, likely from PU, or will not be matched to any track at all. The former case would decrease the object purity. Although the Tk-Electron identification imposes a track p_T cut of 10 GeV as a proxy for the track fit quality, combined with additional isolation requirements to reduce hadronic activity from PU, it still contributes to background rates. Furthermore, in the latter case the true electron will not be reconstructed as a track-matched electron, which decreases the efficiency. To recover this efficiency loss, in particular for low p_T electrons, the geometrical track-cluster matching at the calorimeter surface for the Composite ID is performed with a much looser matching window compared to the tight elliptical matching window of the baseline identification. Based on Figure 7.8 a matching requirement of $\Delta R < 0.2$ is chosen, which is almost fully efficient, even for the softest electrons.

Relaxing the track p_T and matching requirements will increase the electron efficiency, but also increases object multiplicities and background rates. To select only genuine electrons among the composite candidates, a multivariate classification algorithm is developed using both track and cluster based features as input. A BDT based model is chosen for this task, due to its generally lower resource requirements compared to neural networks. This is the third difference with respect

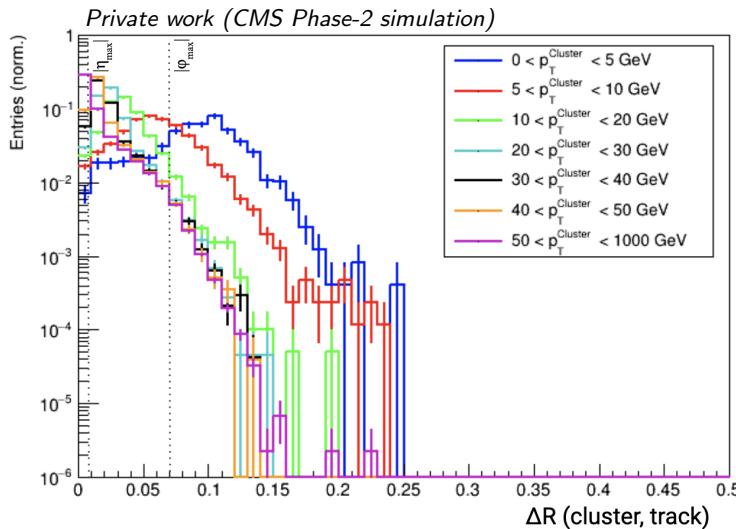


Figure 7.8: Spatial separation between tracks and the highest p_T HGCAL cluster that is matched to a generator level electron. The distributions, binned in cluster p_T , are for single electron events generated at 0 PU. The dotted lines at $\Delta R = 0.0075$ and 0.07 correspond to $|\eta_{\max}|$ and $|\phi_{\max}|$, respectively, which define the elliptical cut used for the baseline identification.

7.3 Improving Performance for Track-matched Electrons

to the baseline identification: the baseline identification relies on track information to match a cluster *after* the cluster is already identified using the tight working point of the stand-alone e/γ ID. The Composite ID instead exploits the track information as input to the model, not for a-posteriori matching. As described in more detail in Section 7.3.2, the model will make use of three feature classes; Besides track and cluster features, also several *matching features* are defined based on kinematics of both object types, such as $\Delta\eta(\text{cluster},\text{track})$. With these different feature classes, the model is able to simultaneously control the individual object qualities *and* the tightness of matching. This is the crucial benefit of the Composite ID approach, which allows both high and low p_T electrons to be identified in a more optimal way, despite their different appearance due to e.g. bremsstrahlung effects.

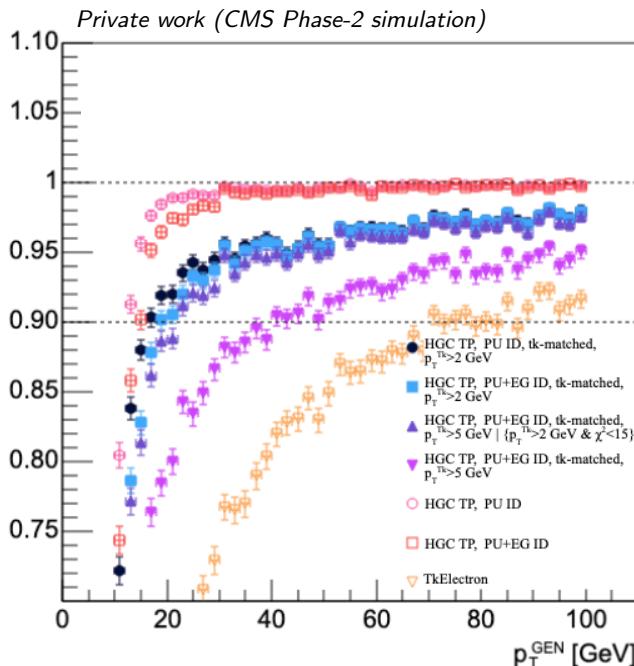


Figure 7.9: Identification efficiency of electrons in the CMS end-caps ($1.52 < |\eta| < 2.4$) generated with an average of 200 PU events, factorizing the effect of the various pre-selection and identification steps. *HGC TP* refers to HGCAL clusters without any pre-selection requirements. *Tk-matched* refers to a cluster that is matched to a L1 track via the loose $\Delta R < 0.2$ cut. *EG ID* refers to clusters passing the loose WP of the stand-alone e/γ ID. The maximum efficiency that can be achieved with the Composite ID is shown with the purple triangles.

With the Composite ID it is possible to achieve a major increase in efficiency with respect to the baseline identification. Figure 7.9 shows the efficiency of composite candidates after the pre-selection but before the multivariate classification, as well as the efficiency for the baseline TkElectron ID. The difference corresponds to the maximum efficiency that can be recovered with the Composite ID. It can be seen that for electrons with $p_T > 50$ GeV up to 10% efficiency can be recovered. Furthermore, the maximum efficiency for electrons of 10 GeV is roughly 70%, in contrast to the negligible efficiency of the baseline algorithm. Even though this does not include the performance of the multivariate classification of the Composite ID, and the efficiencies do not correspond to a fixed rate, the maximum efficiency that can be recovered by using composite candidates does show great potential of this new identification strategy.

7.3.2 Model Training

This section and Section 7.3.3 describes the Composite ID model training and performance in MC simulations. The model was initially developed with the goal of maximizing the performance and establishing a benchmark before simplifying it to fit within the hardware constraints. To also satisfy the hardware constraints, minor changes to the training procedure were necessary. These, and the effect on the model performance, will be described in Section 7.3.4.

The Composite ID model is implemented with a BDT that was trained as binary classifier, labeling composite candidates with a probability to be a genuine electron or a background candidate, typically from PU. Even though more complex machine learning techniques, such as neural networks, could also be used, BDTs are preferred due to their lower resource requirements and efficient implementation on FPGAs. To train the model, a large sample of signal and background composite candidates was constructed from the following simulated data sets:

- **Signal:** DOUBLEELECTRON_PT1TO100, containing final state electrons from the hard interaction process with flat p_T distribution between 1 and 100 GeV, simulated with an average of 200 PU events.
- **Background:** MINBIAS, containing only particles from soft proton-proton collisions, simulated with on average 200 such PU events.

In both data sets, composite candidates are constructed from any pre-selected L1 track that is geometrically matched to any pre-selected HGCAL cluster via $\Delta R(\text{cluster}, \text{track}) < 0.2$. For the track-cluster matching, the track is propagated

7.3 Improving Performance for Track-matched Electrons

to the calorimeter surface to account for the effect of the magnetic field. Signal candidates are taken from the DOUBLEELECTRON sample, requiring that the composite candidates are matched to a generator level (GEN) electron. To match to a GEN electron, the composite candidate must pass $\Delta R(\text{GEN, cluster}) < 0.2$, again using the coordinates of the objects at the calorimeter surface. In order to improve the purity of the signal sample, the candidate must additionally pass $\Delta z(\text{GEN, track}) < 1 \text{ cm}$, using the coordinates at the interaction vertex. All these matching requirements are illustrated in Figure 7.10. Background candidates are taken from the MINBIAS sample, and are not required to be matched to a generator level electron. As such, the background sample consists of a mixture of candidates from soft proton interactions³.

The model is trained using features that can be divided into three classes: cluster features, track features, and matching features. The cluster and track features can be used to control the quality of the respective objects, while the

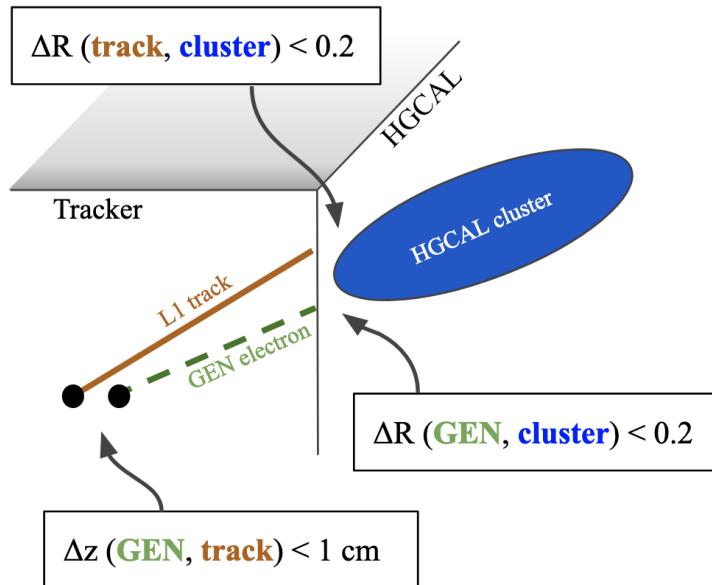


Figure 7.10: Schematic of the Tracker-HGCAL interface in the r - z plane, showing the object matching requirements used to construct composite candidates and define the signal sample for the model training.

³The mixture in principle does also include a small fraction of genuine electrons from pileup collisions, which are regardless considered as background.

matching features - defined from the kinematics of both objects - give control over the tightness of the object matching. All features and their definitions are shown in Appendix C. Below follows a summary.

- **4 track features** are used, including the p_T and z-coordinate of the track, as well as the quality (χ^2) of the track fit. Also the number of stubs used to construct the track is added.
- **20 cluster features** together give a detailed description of the particle shower in HGCAL. Typical features are the number of layers with energy deposits, coordinates of the energy-weighted center of the cluster, and the ratio of energy coming from hadronic decays with respect to the total cluster energy. Also several quantile related features are defined to describe the penetration depth of the particle in HGCAL.
- **3 matching features** are defined using the p_T , η and ϕ coordinates of the cluster and the track.

It should be noted that at the time of writing, the exact data formats of trigger primitives sent from the Tracker and HGCAL backends to the L1T, as summarized in Tables 7.1 and 7.2, are not yet finalized. Furthermore, the computation of some features can be expensive in terms of latency or hardware requirements, and using all 27 features together may result in a BDT that does not fit within the FPGA resource budget. The feature set used for the final model will therefore be subject to changes, as will be discussed in Section 7.3.4. On the other hand, it is clear that many of the features, in particular those related to the clusters, contain similar information. This can be seen by the high correlation between cluster features, shown in Figure 7.11. Removing redundant features from the training can therefore be done without loss of performance. Lastly, besides skimming down the feature set, also new variables are being explored. Track quality MVAs evaluated in the GTT or quantities sensitive to bremsstrahlung could be valuable input for the Composite ID and will have to be tested in future. The signal and background distributions of a selection of input features is shown in Figure 7.12.

The discrimination power of the input features depends strongly on the p_T of the composite candidates; In particular, features that are correlated with the electron p_T will have obviously different distributions for energetic candidates from the hard scatter and for soft candidates from PU interactions. Analogously to the stand-alone e/γ ID that was trained in two $|\eta|$ bins [148], it was found that training the model independently in two p_T bins resulted in better background rejection in

7.3 Improving Performance for Track-matched Electrons

comparison to training a single discriminator inclusively in p_T . The two p_T bins are defined as:

- Low p_T bin: $5 < p_T^{\text{cluster}} < 25 \text{ GeV}$
- High p_T bin: $p_T^{\text{cluster}} > 15 \text{ GeV}$

This way, the low p_T and high p_T models may rely differently on the input features in regions that are respectively very rich and sparsely populated by PU candidates, hereby achieving more optimal discrimination power. The boundaries are furthermore chosen to have an overlapping margin at $15 < p_T < 25 \text{ GeV}$. As such, for the implementation of these models in the various electron or cross-object triggers, a combination of both models could be used, with the choice of model evaluation depending on the candidate p_T and working points defined to obtain

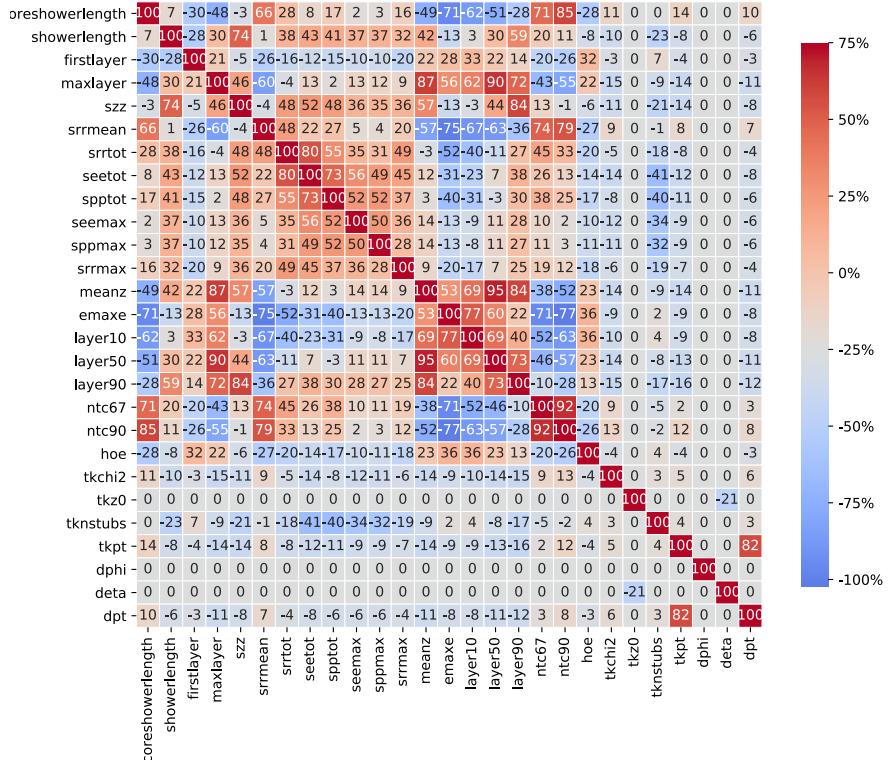


Figure 7.11: Feature correlation matrix corresponding to the high p_T training of the Composite ID. Feature definitions are shown in Appendix C.

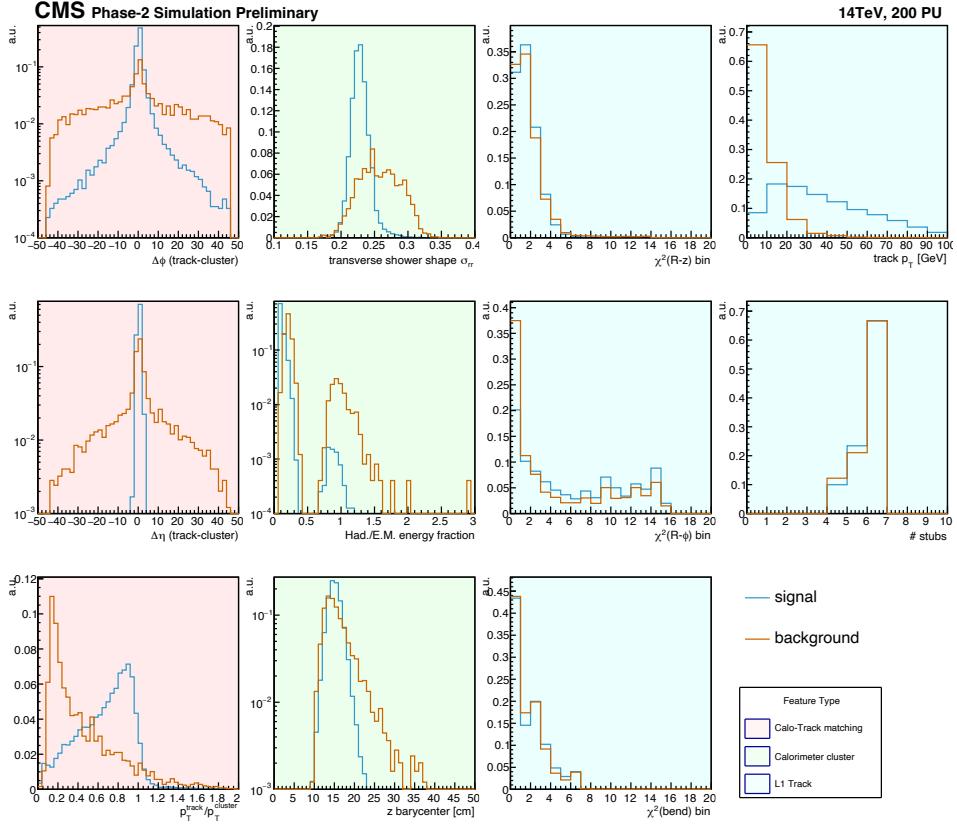


Figure 7.12: Normalized signal and background distributions of input features of simulated composite objects used in the Composite ID model training [244].

smooth background rejection over the full spectrum.

The Composite ID BDTs were trained with XGBoost(v1.5.0), an open-source software package which provides a regularizing gradient boosting framework for Python [248]. The hyper-parameters were configured to construct a model of 10 booster trees, each with a maximum depth of 6 layers. For the firmware implementation, described in Section 7.3.4, the maximum tree depth was reduced to 4 in order to decrease resource requirements.

7.3.3 Performance in simulation

After training BDTs, they are typically characterized by their *receiver operating characteristic* (ROC) curves, which display signal efficiency as function of back-

7.3 Improving Performance for Track-matched Electrons

ground rejection. The ROC curves for the low p_T and high p_T trainings of the Composite ID are shown in Figure 7.13 and compared to those of the stand-alone e/γ ID that is used for the TkElectron. The efficiencies are computed using composite candidates divided according to the p_T binning of the Composite ID training. Since the stand-alone e/γ ID was trained on clusters with $p_T > 20$ GeV, part of the major improvement of the Composite ID shown in the left plot of Figure 7.13 is attributed to the use of the stand-alone ID outside its training range. That said, the right plot clearly proves that the addition of Tracker information in the multi-variate ID contributes to a notably better discriminator.

From these ROC curves working points (WP) are defined that correspond to a fixed efficiency and background rejection inclusively in the respective p_T bins. In the remainder of this chapter, WPs are labelled with the p_T bin of the Composite ID training (high or low) and the signal efficiency. For example, “WP90% of the low p_T model” corresponds to a BDT score threshold that gives 90% signal efficiency with the low p_T training, based on composite candidates with $5 < p_T < 25$ GeV. The corresponding level of background rejection can be read from Figure 7.13. Furthermore, it should be noted that these conventions are merely WP definitions. To evaluate the performance of a certain model or WP over the full p_T spectrum, the candidates are not binned in p_T ; Even though the models were trained in p_T bins, they can still be used to identify electrons with

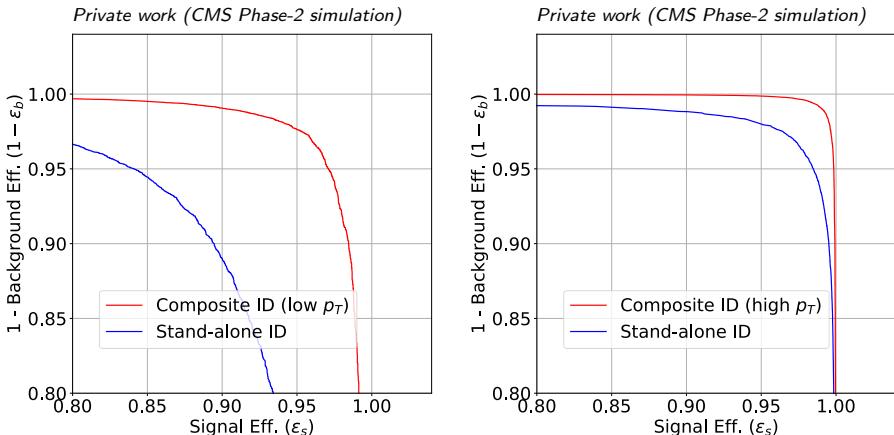


Figure 7.13: ROC curves of the Composite ID (red) computed by evaluating the low p_T training (left) and high p_T training (right) models on composite candidates divided according to the respective p_T binnings. The stand-alone ID (blue) refers to the cluster-only ID that is used as part of the baseline TkElectron identification.

p_T outside the training ranges.

After defining a set of WPs, the physics performance is evaluated using the identification efficiency as function of the p_T of generator level electrons, and the trigger rate as function of the Level-1 reconstructed (cluster) p_T threshold. These are shown in Figure 7.14. A first thing to note is that the various WPs of both the low and high p_T models show a distinct performance in terms of efficiency and rate. In other words, the Composite ID allows tuning the overall trigger

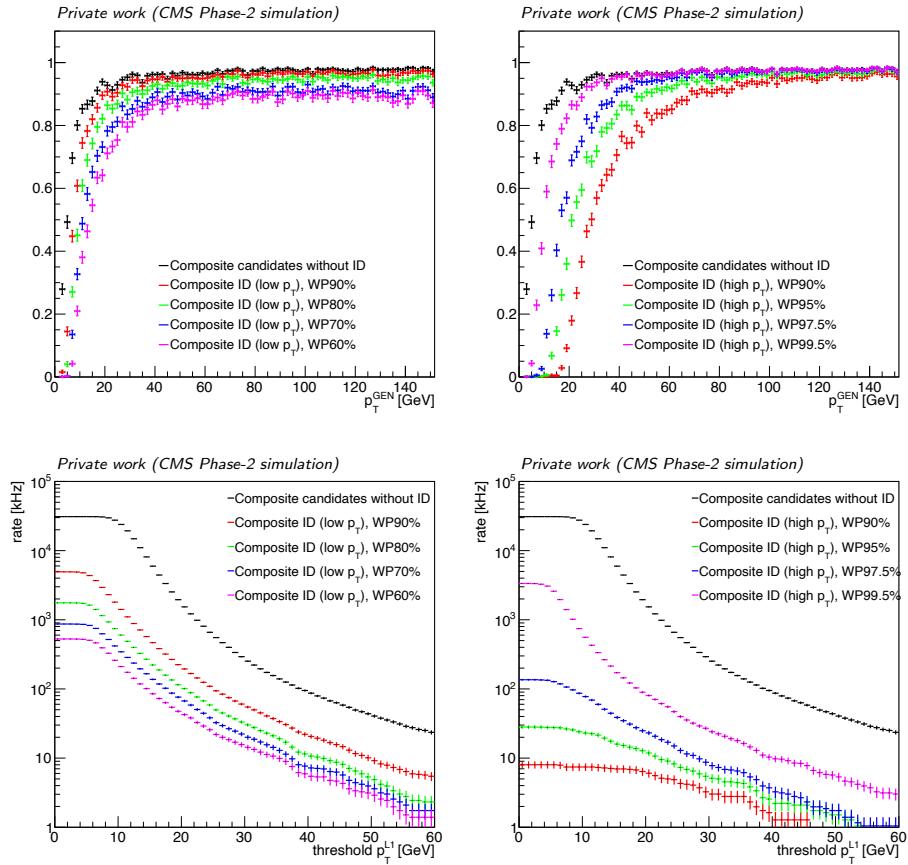


Figure 7.14: Performance of the Composite ID based on simulated electrons in the CMS end-caps ($1.52 < |\eta| < 2.4$). Top: Identification efficiency as function of generator-level electron p_T . Bottom: Single electron trigger rates as function of the Level-1 reconstructed (cluster) p_T threshold. The rates are normalized assuming the LHC orbit frequency of 11.2 kHz and 2800 proton bunches per beam. The left (right) column corresponds to the low (high) p_T trainings of the Composite ID. The efficiencies and rates are computed using DOUBLEELECTRON and MINBIAS samples, respectively, simulated with an average of 200 PU events.

7.3 Improving Performance for Track-matched Electrons

performance simply by picking a different WP, which contributes to a robust and flexible trigger. This is in contrast to the baseline identification, which mainly relies on the track p_T cut of 10 GeV to control the background rates, resulting in reduced flexibility in particular at the low end of the p_T spectrum.

The low and high p_T models show a striking difference between their efficiencies as function of the WPs. For the high p_T training, the model reaches a plateau efficiency of approximately 100% regardless of the WP. Using different WPs affects mostly the efficiency at low p_T (< 40 GeV). In other words, the model has learned that the background from PU populates mostly the low p_T region, which can be reduced by cutting tighter on features that are correlated with the p_T of the track or cluster⁴. The low p_T model, on the other hand, does not use this p_T dependency, since the model was trained using (signal) candidates with similar p_T as the bulk of the PU background. Consequently, using a different WP affects the efficiency over the full p_T spectrum, i.e. also at the efficiency plateau.

With the presence of 200 PU interactions, it is inevitable that some composite candidates will consist of a “true” electron cluster that is matched to a PU track. Even though this may increase the identification efficiency, it also contributes to a decreased purity that may result in much higher rates. Figure 7.15 shows the p_T response and resolution along the z-axis of the L1 track - part of a composite

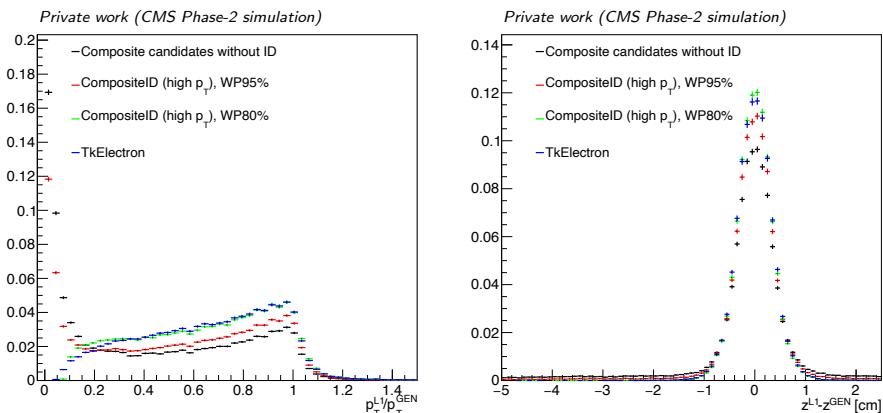


Figure 7.15: Left: Track p_T response of the Level-1 reconstructed objects with respect to the matched generator level electron in the CMS end-caps ($1.52 < |\eta| < 2.4$). Right: Resolution of the z-coordinate of the track with respect to the generator level electron. Both plots are computed using single electron events simulated with on average 200 PU events.

⁴In practice the nodes of the BDT will give different weights to the final score depending on cuts on p_T -correlated features.

candidate - with respect to a generator level electron, simulated with an average of 200 PU events. The p_T resolution of composite candidates before applying any ID shows a clear peak at 0. Furthermore, the z resolution shows a non-negligible contribution in the tails. Both plots therefore confirm that some of the composite candidates (before applying a cut on the BDT model score) consist of a cluster that is wrongly matched to a low p_T track from a PU object. As mentioned several times, the baseline identification minimizes this contribution from PU tracks by requiring track $p_T > 10$ GeV. For the Composite ID this requirement is relaxed, allowing tracks down to 2 GeV to be used. Even though this introduces many more PU tracks, Figure 7.15 shows that the Composite ID is able to control this contribution down to a level similar to the baseline TkElectron identification, simply by tightening the WP. This capability was observed for both high and low p_T models.

The final physics performance of the Composite ID is compared to baseline identification, by computing the electron identification efficiency at a fixed rate. Based on the rate plots of Figure 7.14, trigger thresholds on the (cluster) p_T are derived that correspond to a rate of 20 kHz for the high p_T model and 150 kHz for the low p_T model. The efficiency curves, including the resulting cuts on cluster p_T , are shown in Figure 7.16. The Composite ID shows a clear improvement with

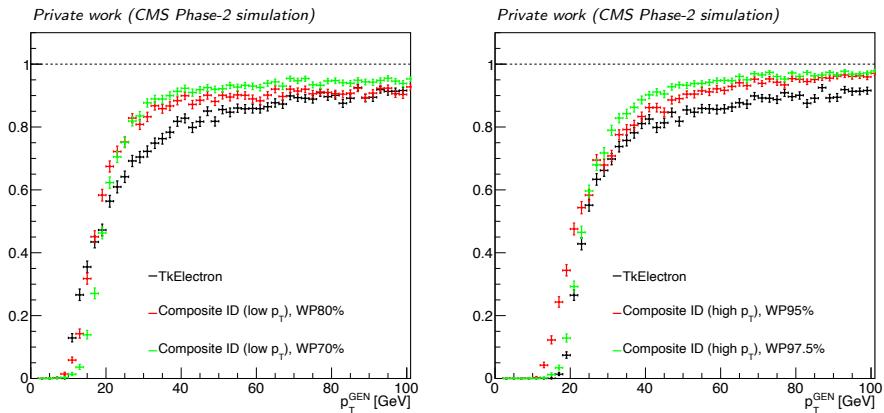


Figure 7.16: Identification efficiency of electrons in the CMS end-caps ($1.52 < |\eta| < 2.4$) as function of generator-level electron p_T , corresponding to a fixed rate of 150 kHz and 20 kHz for the low p_T model (left) and high p_T model (right). The efficiencies of baseline TkElectron identification in both plots are based on the same stand-alone ID discriminator, using the standard tight WP. Only the L1 p_T thresholds differ to match the respective rates. Both plots are computed using single electron events simulated at an average of 200 PU events.

the high p_T training. By picking e.g. WP95%, it is possible to roughly double the efficiency for electrons with $p_T = 20$ GeV, while the plateau efficiency is improved by 5-10%. While these results are fully based on simulation, ignoring implications from running this trigger strategy on FPGAs, it is evident that the Composite ID has a great potential to improve the efficiency with respect to the baseline identification over the full p_T spectrum. For the low p_T training, the results remain inconclusive. Even though the “shoulder” of the efficiency is sharper, the most competitive WPs of the Composite ID show similar turn-on and plateau efficiency to the TkElectron. It should be noted, however, that the trigger thresholds derived for a rate of 150 kHz are relatively high with respect to the p_T range of the low p_T training. By evaluating the model for candidates with lower p_T , improvements with respect to the TkElectron might become more evident. This could be achieved by considering e.g. double-electron or cross-object triggers, such as $ee + p_T^{\text{miss}}$ triggers⁵, which typically apply electron p_T thresholds of $O(10)$ GeV. At the time of writing, the use of the low p_T model in dielectron or cross-object triggers is still being investigated. Therefore, the remainder of this chapter will focus on the firmware implementation of the high p_T model only.

7.3.4 Firmware implementation

While the results shown in the previous section promise a clear improvement of efficiency at fixed rate with the Composite ID compared to the baseline identification, it should be noted that the model was not optimized for deployment in firmware. The Composite ID will have to be implemented in the Correlator Trigger boards, running on limited FGPA resources. These resources are used not only by e/γ algorithms, but also by eg. PF and PUPPI. Furthermore, the presence of multiple composite candidates per event, combined with the latency constraints, implies that multiple instances of the model need to be evaluated in parallel. It is therefore imperative keep the resources that are required to evaluate the Composite ID BDT under control. Moreover, while the results of the previous section are based on simulated objects and floating point precisions, in practice the model will have to be evaluated in hardware using limited bit-precision for the representations of the composite candidates. The number of bits is kept to a minimum in order to

⁵In Phase 1, the muon variant of such triggers ($\mu\mu + p_T^{\text{miss}}$) was already used: it is in fact one of the main triggers for the search presented in Chapter 5! For $ee + p_T^{\text{miss}}$ triggers the rate in Phase 1 was found to be unsustainable. However, the Composite ID – which exploits the new Tracker tracks and machine learning at L1 – may enable the use of $ee + p_T^{\text{miss}}$ triggers in Phase 2, which would be of great value particularly for new physics searches.

satisfy the constraints on the bandwidth of the L1 infrastructure. Besides estimating the required FPGA resources, the algorithm therefore also has to be emulated with quantized features, which could reduce the final physics performance.

Recent developments have lead to the introduction of a new software package, known as CONIFER [246], which allows BDTs to be translated to FPGA firmware for low latency inference. Furthermore, it can be used to estimate the resource and latency requirements for the BDT inference, while allowing for flexible clock frequency, bit precisions and target FPGA chips.

As was already hinted in the previous section, it is clear that the (high pt) algorithm relies on a sizable classifier that consists of 10 booster trees, each with a maximum of 6 layers per tree, and 27 features to identify composite candidates. Synthesizing the BDTs to the FPGA therefore resulted in a model that was too large for the given firmware constraints. Furthermore, it was found that the calculation of several track and cluster features of the TPs was too expensive in terms of the required computational resources on the backend electronics. Those features will therefore not be available in the L1T firmware and hence cannot be used by the Composite ID. With the aim of preserving maximum discrimination power, the Composite ID was therefore retrained with minor modifications, in order to reduce the model complexity and rely only on features that will be available in firmware.

Firstly, the number of input features was reduced from 27 to 11, using only those that can be implemented in the backends that generate the TPs. The features were selected based on their relative importance to the final prediction. Typical

Table 7.4: Bit representations of the selected input features of the Composite ID. For features of type *ap_ufixed*, the difference between the two numbers defines the number of bits representing the value below the decimal point, while the second number is used to define the bits representing numbers above the decimal point. For all other types, the total number of bits is stated.

Type	Feature	Precision in firmware
Cluster	HOE	ap_ufixed <10,5>
	MEANZ	ap_uint <8>
	SRRTOT	ap_ufixed <10,1>
Track	TKPT	ap_ufixed <14,12>
	TKCHI2RZ	ap_uint <4>
	TKCHI2RPHI	ap_uint <4>
	NSTUBS	ap_uint <3>
Matching	DETA	ap_int <10>
	DPHI	ap_int <10>
	DPT	ap_fixed <21,12>

7.3 Improving Performance for Track-matched Electrons

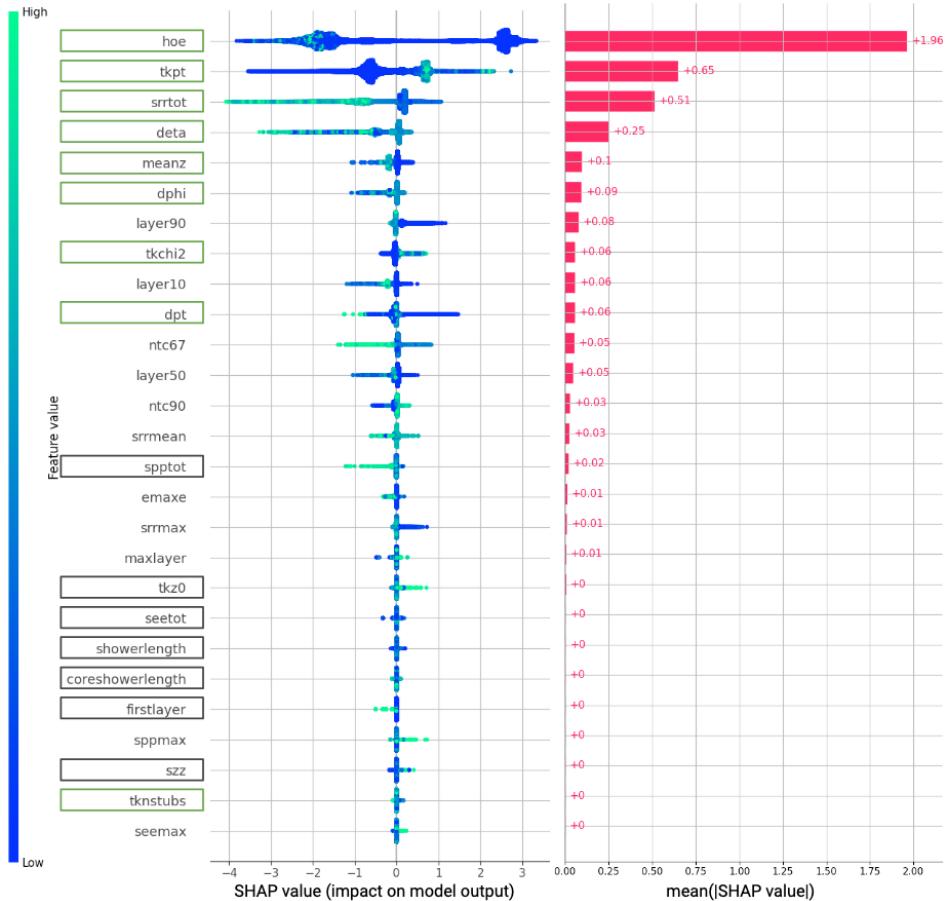


Figure 7.17: SHAP values per feature per candidate used to train the Composite ID. Features are displayed in descending order of relative importance and marked with a (black or green) box if the feature can be computed by the TPG backends. Only the features in green boxes are selected for the model retraining. It should be noted that the `TKCHI2` variable in firmware has since been redefined into three components (`TKCHI2RZ`, `TKCHI2RPHI` and `TKCHI2BEND`), which are shown in Figure 7.12 and are all used for the model retraining. The definition of each variable is shown in Appendix C.

figures of merit to quantify this importance is the number of times a feature is used to split a branch across all trees, or the average training loss reduction gained when using a certain feature. However, for both these measures of feature importance, higher scores do not always mean the model relies more on those features to make the prediction. A new figure of merit was recently proposed, based on Shapley Additive Explanations (SHAP) values, to provide a better consistency with human intuition, and allow one to explain individual predictions based on the feature values of the input candidate [249]. Figure 7.17 shows the SHAP value per feature for each candidate used to train the Composite ID. The higher the $|SHAP|$ value, the higher the impact of a given feature on the model prediction. The features are sorted by the mean of the $|SHAP|$ values, hence the feature importance is displayed in descending order. It can therefore be seen that the fraction of hadronic energy of the cluster (HOE) and the track p_T (TKPT) are the two most important features to identify electrons with the Composite ID. Based on this ranking, and on the firmware compatibility of the features, 11 features are selected: the 3 highest scoring cluster features, 3 matching features and 5 track features. The selected features are highlighted with the green boxes in Figure 7.17 and the corresponding distributions of signal and background candidates used for training are shown in Figure 7.12.

A second modification to the model training involved reducing the maximum tree depth. The FPGA resources needed to evaluate a BDT strongly depend on this hyper parameter, with in particular the number of look-up tables (LUTs) increasing exponentially with the maximum tree depth [246]. It was found that a maximum depth of 4 was the optimal trade-off between identification performance and model complexity (i.e. required resources).

The third and last modification to the training procedure was using quantized representations of the 11 selected input features, instead of floating point precision. In the L1T firmware only a small number of bits will be reserved to store each feature of the clusters and tracks. Even though the exact bit representations are still to be finalized, the latest representations are shown in Table 7.4. This choice of representation for each feature is based on a delicate trade-off between feature resolution, its importance in making physics decisions, and bandwidth constraints. To train the final model, the samples of signal and background candidates were fully emulated according to these bit representations.

After reducing and quantizing the input features, and decreasing the maximum tree depth, the Composite ID was retrained using the same procedure as described in Section 7.3.2. The resulting signal and background efficiencies are

7.3 Improving Performance for Track-matched Electrons

shown in Figure 7.18 (left plot, red and blue lines). Next, the retrained model was implemented in the L1T emulator and firmware; While the firmware implementation defines the bit-wise logic operations performed on the physical Correlator Layer-1 trigger boards, the C++ based emulator of the L1T serves as a tool to cross-check the firmware and test the expected physics performance using simulation, with precisions capped to that of the firmware. The translation of the XGBOOST model to FPGA firmware and a version that can be run in the C++ based emulator was performed with the CONIFER package. Since CONIFER allows for only one bit precision to represent all features, the $ap_fixed<21,12>$ representation was used for the model synthesis of both the firmware and emulator versions, as it accommodates the full precision of each of the features in the firmware (see Table 7.4).

The BDT inference with the retrained XGBOOST model and the C++ based emulator model is compared in Figure 7.18. Even when given the same composite candidates, non-trivial differences in the BDT inferences may occur, since the computations leading to the output score in the emulator (and firmware) are bit-precise, while those in the XGBOOST model rely on floating point precision – regardless of the quantization of the input features themselves. However, the figure shows that an excellent agreement is achieved and the impact of the quantization of the BDT inference is small.

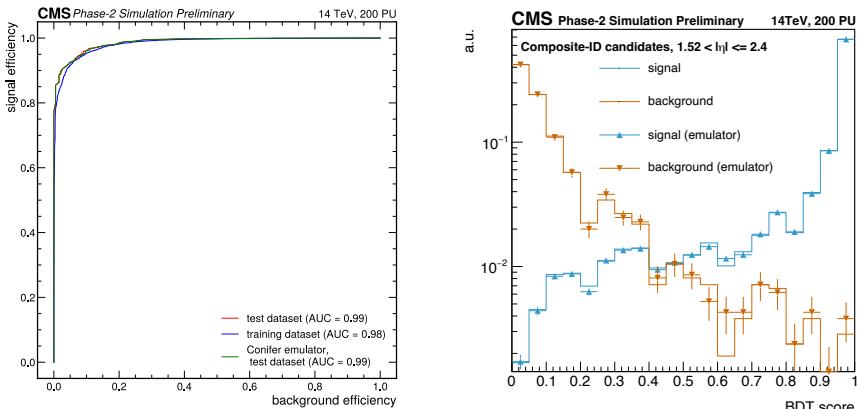


Figure 7.18: Left: ROC curves for the XGBOOST and emulator versions of the Composite ID BDT. Both versions rely on quantized input features, but only the emulator version uses bit-precise computations for the output score. Right: BDT output scores shown for signal and background candidates using both versions of the Composite ID BDT. [244]

In firmware, the Composite ID was implemented as a functional block of logic, referred to as an *intellectual property* (IP) core, in the Correlator Trigger Layer-1 (CTL1) firmware design using the Vitis HLS synthesis tools and the CONIFER package. The number of parallel BDT copies that are needed to process all candidates depends on the frequency of data transfers to the CTL1 and the frequency of the e/γ IP core. The CTL1 boards serving the CMS end-cap regions have a time-multiplexing (TMUX) factor of 6, meaning that each board has 6 bunch-crossings to accept data from a full end-cap before the data from the next event come in. With data transmission at a clock-frequency of 360 MHz and a collision rate of 40 MHz, one bunch-crossing is equivalent to 9 clock-cycles. As such, the CTL1 boards have 54 clock-cycles before the next event comes in. The end-cap data is divided into 9 regions, such that the e/γ core instance on a given board accepts one region every 6 clock-cycles. This would imply the initiation interval (II) for the e/γ algorithms equals 6. However, the e/γ algorithms instead operate with a clock-frequency of 180 MHz and an II of 3. The e/γ IP core will therefore need to accept data from a new region (or event) every 3 clock-cycles (at 180 MHz). To decide how many parallel copies of the Composite ID are needed, the maximum number of HGCAL clusters per region per event, as well as the maximum number of tracks per cluster (assuming the matching procedure of the Composite ID) are shown in Figure 7.19. Based on this figure, and with sufficient safety margin – the HGCAL TP generation is not yet fixed – an upper bound of 10 is set on the

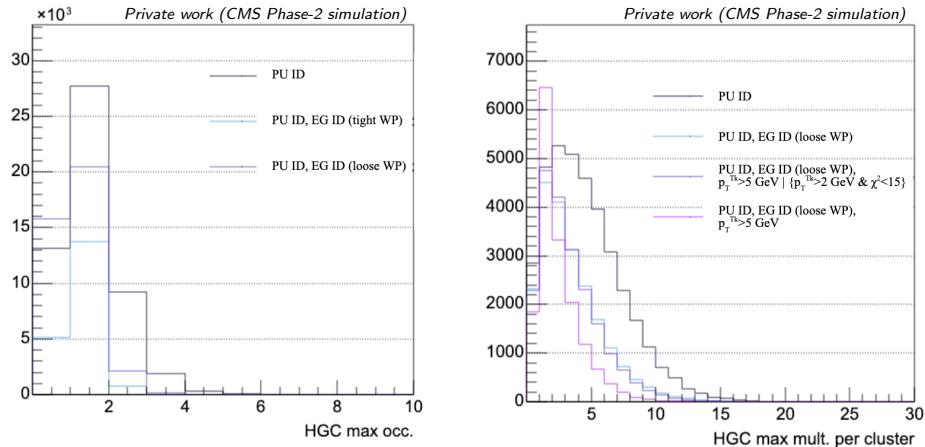


Figure 7.19: Left: The maximum number of HGCAL clusters per PF region. Right: The number of tracks that are matched to a given cluster. Both plots are based on $t\bar{t}$ events, simulated with on average 200 PU events.

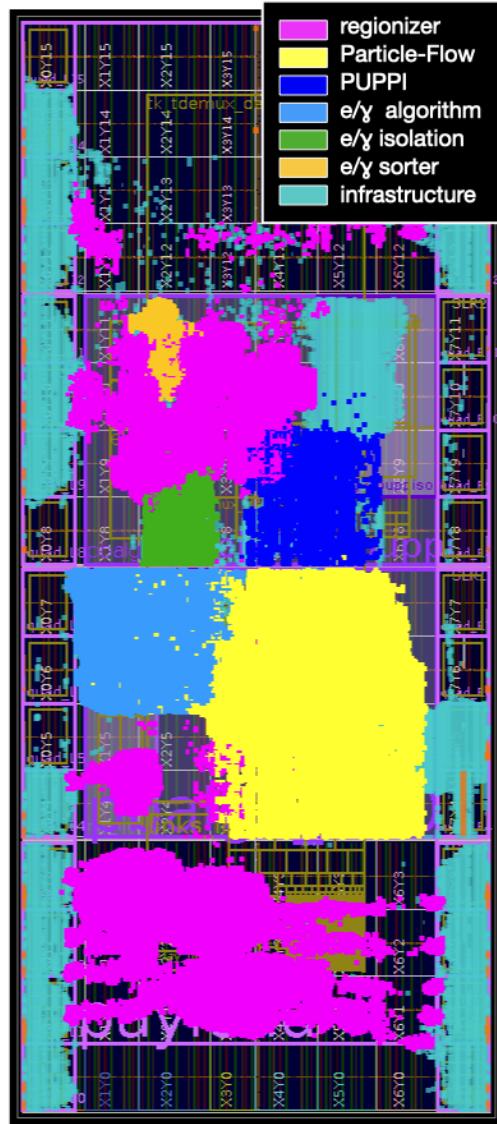


Figure 7.20: Floorplan of a Xilinx Vertex UltraScale+ VU13P FPGA running the Correlator Layer-1 trigger algorithms. The various colors correspond to the different IP cores of the firmware logic. The Composite ID is evaluated in the e/γ algorithm shown in light blue.

	LUT	FF	BRAM	DSP
e/γ IP	3.1%	0.4%	0.0%	1.6%
Total	24.4%	17.6%	29.5%	14.3%

Table 7.5: Fractions of the main resources of the Xilinx Virtex UltraScale+ VU13P FPGA used by the e/γ IP core and the overall end-cap Correlator Layer-1 firmware. For a clock frequency of 180 MHz and an initiation interval of 3, the latency is 18 clock cycles.

number of clusters per region. Furthermore, an upper bound of 3 is set on the number of tracks (per cluster) that are used to construct a composite candidate. If a cluster matches to more than 3 tracks, only the 3 best track-cluster matches are used, which minimize the difference in object p_T . Even though this choice (which will be optimized in future) introduces some inefficiency, it limits the required FPGA resources by reduces the maximum amount of inferences to roughly 10 (clusters per event per region) \times 3 (tracks per cluster) = 30 candidates per event per regions. With an II of 3, this means that 10 parallel copies of the model are needed.

The full Composite ID algorithm implementation in firmware was demonstrated running on a Xilinx Virtex UltraScale+ VU13P FPGA (which has roughly 8 times more resources than the FPGAs used in the Phase 1 L1T). Table 7.5 shows the main FPGA resources required by the e/γ IP core that includes the Composite ID, as well as those required by the overall end-cap Correlator Layer-1 firmware. Figure 7.20 shows the corresponding FPGA floorplan. As can be seen, the Composite ID fits well within the available resources.

The Composite ID algorithm in firmware was tested by running the firmware

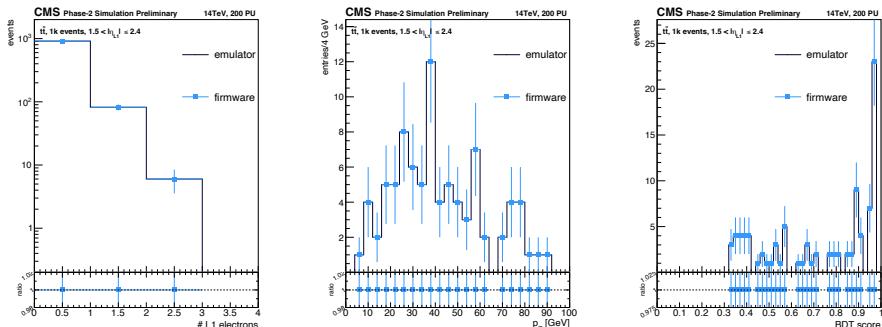


Figure 7.21: Comparison of the number of L1 electrons, electron p_T and Composite ID BDT score, when running the Composite ID algorithm in the L1T firmware and the emulator on 1000 $t\bar{t}$ events [244].

and emulator on the same 1k $t\bar{t}$ events and comparing event- and object-level distributions, as shown in Figure 7.21. Bit-wise agreement was observed for all variables, including the BDT output score. Consequently, the final physics performance may be evaluated similarly as in Section 7.3.3, but with the fully emulated algorithm. The trigger rates of the Composite ID are shown as function of the Level-1 reconstructed (cluster) p_T threshold in Figure 7.22 (left), and compared to those of the baseline identification algorithm. The p_T thresholds are derived for single electron triggers allowing 9 kHz of rate and the corresponding turn-on curves are shown in Figure 7.22 (right). For the chosen WP of the Composite ID, a clear improvement with respect to the baseline identification is observed over the full p_T spectrum: an earlier turn-on efficiency is achieved at low p_T and the Composite ID recovers 5-10% efficiency for electrons with $p_T > 30$ GeV.

7.4 Conclusion

With the Composite ID a new electron identification strategy is introduced for the Phase 2 L1T, replacing the baseline identification presented in Ref. [148]. A minimum of 5-10% efficiency for electrons with $p_T > 30$ GeV can be recovered with respect to the baseline identification, using the high p_T training of the Composite

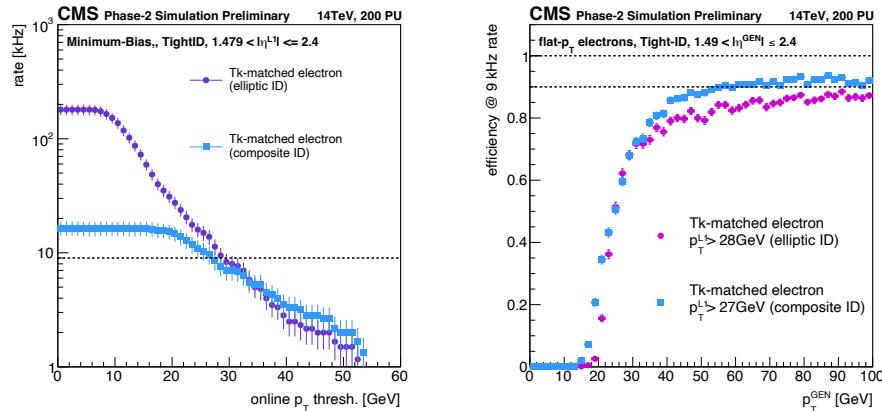


Figure 7.22: Comparison of the physics performance of the emulated Composite ID (blue) and baseline TkElectron identification (purple) [244]. Left: Single electron trigger rates as function of the Level-1 reconstructed (cluster) p_T threshold, normalized assuming the LHC orbit frequency of 11.2 kHz and 2800 proton bunches per beam, and computed with MINBIAS samples with on average 200 PU events. Right: Identification efficiency as function of generator-level electron p_T corresponding to a fixed rate of 9 kHz, computed using DOUBLEELECTRON samples with on average 200 PU events.

ID. The Composite ID is successfully implemented in the VU13P FPGA of the Correlator Layer-1 Trigger boards and has achieved 100% bit-wise agreement with the emulator used to quantify the physics performance. As such, the Composite ID is the first machine learning algorithm within the CMS Phase 2 L1T project that is fully demonstrated in the L1T emulator and firmware, and hereby paves the way for other trigger strategies based on machine learning techniques. Studies are still ongoing to explore the potential of the low p_T training to be used in cross-object triggers, such as $ee + p_T^{\text{miss}}$ triggers, which generally apply lower p_T thresholds than the single electron triggers shown in this section.

Summary

Theories beyond the Standard Model predict the existence of new physics at the electroweak scale in order to provide answers to several fundamental questions related to eg. gravity, the hierarchy problem and the nature of Dark Matter. This thesis aids in the continued exploration of the *high energy* electroweak scale, making use of *low energy* leptons as a probe for new physics.

The search presented in Chapter 5 targets new physics realizations with compressed mass-spectra, which are motivated by eg. models of Supersymmetry, and characterized by only low energy (soft) Standard Model particles from cascade decays. Signal scenarios with smaller mass-splittings (Δm) become increasingly challenging to probe, due to low signal acceptance, decreased object reconstruction efficiencies and the overwhelming hadronic background of the LHC environment. As such, these scenarios could have escaped experimental observation thus far. This search targets final states with two or three soft leptons and a moderate amount of missing transverse momentum (p_T^{miss}), making use of 138 fb^{-1} of proton-proton collisions at a center of mass energy $\sqrt{s} = 13 \text{ TeV}$, collected by the CMS experiment during LHC Run 2 (2016-2018). The sensitivity to more compressed signal scenarios is greatly increased by the usage of a new electron reconstruction procedure, that facilitates transverse momentum (p_T) thresholds as low as 1 GeV. The background from Standard Model processes consists primarily of non-prompt and fake leptons associated to hadronic activity. These backgrounds are reduced via tight object and event selection criteria and estimated with dedicated data-driven techniques. The analysis furthermore employs a new search category based on the presence of soft *and* displaced muon pairs to probe – for the first time in CMS – both mass-compressed *and* long-lived new particle states. A key kinematic variable in this search is the invariant mass ($m_{\ell\ell}$) of oppositely charged pairs of leptons, which serves as a proxy of the mass-splitting for signal events and pro-

Summary

vides a great handle to distinguish background events. To maximally exploit its discrimination power, a novel signal extraction procedure is designed, that consists of optimized $m_{\ell\ell}$ binnings of the search regions for each signal hypothesis. The results of the search are interpreted in terms of a simplified wino-bino model of R-parity conserving Supersymmetry. The search sets exclusion limits at 95% confidence level as function of the lifetime and masses of the new particle states. For promptly decaying winos, the signal exclusion extends to mass-splittings as low as 900 MeV for a wino mass of 100 GeV. This result hereby closes the sensitivity gap at mass-splittings between 0.9 and 2 GeV, implying that LHC searches now exceed the mass-bounds for compressed electroweakinos from LEP experiments over the full range of the mass-splitting. For long-lived scenarios the maximum exclusion is obtained for a $\tilde{\chi}_2^0$ lifetime of $c\tau = 100$ mm, $m_{\tilde{\chi}_2^0} = 400$ GeV and a mass-splitting $\Delta(m_{\tilde{\chi}_2^0}, m_{\tilde{\chi}_1^0}) = 20$ GeV.

Chapter 6 reports the legacy Run 2 combination of searches for electroweak Supersymmetry, that aims to maximally exploit the data-set collected by the CMS experiment thus far. The selected search channels target a variety of (semi-)leptonic and hadronic signatures, and include the search of Chapter 5 as well. In the absence of observed deviations from the Standard Model expectation, limits are set on the pair-production of charginos, neutralinos and sleptons, based on simplified models of R-parity conserving Supersymmetry. The combined search increases the excluded sparticle masses by up to 125 GeV with respect to the most sensitive component searches. Furthermore, for the first time in CMS, interpretations are made for scenarios of pair-produced sleptons in the mass-compressed parameter space by exploiting soft lepton final states; Exclusion limits reach up to slepton masses of 215 GeV for a mass-splitting of 5 GeV.

To further improve the analysis sensitivities, more collision data is needed. From 2026 onwards, the LHC will therefore be upgraded to the High Luminosity LHC, which is designed to deliver proton-proton collisions at 7 times the nominal LHC luminosity. Although this would yield enough data to significantly shrink the statistical uncertainties for new physics searches, major upgrades to the CMS detector are required to continue operations and maintain the physics acceptance of Run 2, even in these harsh running conditions of the High Luminosity LHC. The Level-1 Trigger, responsible for the real-time event selection, has therefore been completely redesigned, and novel selection algorithms are being developed to maximize its physics acceptance. In Chapter 7 a new algorithm is presented for the real-time identification of electrons in the CMS end-caps. The *composite electron identification* strategy relies on the new Tracker capabilities - tracks will

be reconstructed for the first time at the collision rate of 40 MHz - as well as the rich and detailed descriptions of particle showers from the new High Granularity Calorimeter. Electron candidates are constructed by geometrically matching tracks to energy deposits in the calorimeter, and identified using a multivariate analysis based on boosted decision trees. Due to recent gains of FPGA processing power and newly developed compiler packages, this machine learning based model can be evaluated within the time and hardware constraints of the upgraded Level-1 Trigger. The composite electron identification algorithm improves the efficiency for electrons with $p_T > 30$ GeV by a minimum of 5-10% with respect to the baseline strategy that was considered thus far. Furthermore, studies with re-trained versions of the composite identification model are ongoing to increase the efficiency also for electrons with lower p_T , which could be used in cross-object (eg. $ee + p_T^{\text{miss}}$) triggers. Searches for new physics realizations, such as those presented in this thesis, would greatly benefit from such triggers in the future era of the High Luminosity LHC.

Appendix A

Limit Setting in Searches for New Physics at the LHC

Many analyses of experiments at the LHC search for new physics phenomena. In the absence of significant deviations between the observed data and the background expectations, limits are typically set on the properties of some signal model, such as the production cross-section or mass of new physics states. This appendix is a non-exhaustive summary of the main statistical components used in limit setting at the LHC, and is mostly based on Refs. [240, 250].

The method commonly employed for computing exclusion limits is based on the modified frequentist approach, often referred to as CL_s , and an asymptotic approximation. In this method, s denotes the expected signal yields and b the expected background yields. A signal strength modifier μ is introduced that basically scales the signal cross-section and/or branching fraction. Both s and b are subject to uncertainties, represented by nuisance parameters θ , such that $s = s(\theta)$ and $b = b(\theta)$.

The first step in providing exclusion limits is to construct the likelihood function:

$$L(\mu, \theta) = P(x|\mu, \theta) \tag{1}$$

$$= \text{Poisson}(x|\mu \cdot s(\theta) + b(\theta)) \cdot p(\tilde{\theta}|\theta) \tag{2}$$

In the case of a binned analysis, the Poisson term amounts to the product of Poisson probabilities to observe n_i events in bins i : $\prod_i = \frac{(\mu s_i + b_i)^{n_i}}{n_i!} e^{-\mu s_i - b_i}$. The term $p(\tilde{\theta}|\theta)$ corresponds to the probability density functions of the systematic uncertainties, with $\tilde{\theta}$ denoting the default value of the nuisance parameters. It parametrizes the lack of knowledge of what the true value of θ might be. A test statistic t_μ is then defined as a measure of how (in)consistent the background-only

Limit setting

A

hypothesis is with the observed data. In the case of limit setting it is given as:

$$t_\mu = -2 \ln \left(\frac{L(\mu, \hat{\theta}_\mu)}{L(\hat{\mu}, \hat{\theta})} \right), \text{ with } 0 \leq \hat{\mu} \leq \mu \quad (3)$$

where estimators $(\hat{\mu}, \hat{\theta})$ are the set of (μ, θ) that maximize the likelihood and $\hat{\theta}_\mu$ is the maximum likelihood estimator of θ given μ . Equation 3 prevents exclusion of a given value of μ in case of an upward fluctuation of data ($\hat{\mu} > \mu$).

Next, the observed value of the test statistic t_μ^{obs} for the given signal strength modifier is obtained by substituting the actual observed data in Equation 3. Also the values of the nuisance parameters $\hat{\theta}_0^{obs}$ and $\hat{\theta}_\mu^{obs}$ that best describe the observed data are obtained by maximizing the likelihood for the background-only and signal+background hypotheses.

Finally, upper limits are determined by finding the range of μ for which $\text{CL}_s(\mu)$ is lower than a chosen threshold α – by convention typically 5% – with $\text{CL}_s(\mu)$ defined as:

$$\text{CL}_s(\mu) = \frac{\text{CL}_{s+b}(\mu)}{\text{CL}_b(\mu)} = \frac{p_\mu}{1 - p_b} = \frac{P(t_\mu \geq t_\mu^{obs} | \mu, \hat{\theta}_\mu^{obs})}{P(t_0 \geq t_0^{obs} | 0, \hat{\theta}_0^{obs})} \quad (4)$$

$$= \frac{\int_{t_\mu^{obs}}^{\infty} f(t_\mu | \mu, \hat{\theta}_\mu^{obs}) dt_\mu}{\int_{t_0^{obs}}^{\infty} f(t_\mu | 0, \hat{\theta}_0^{obs}) dt_\mu} \quad (5)$$

The functions f in the numerator and denominator are probability density functions of the test statistic, corresponding to the signal+background and background-only hypotheses, respectively. In a typical LHC analysis, these are not trivial to obtain analytically, but can instead be obtained numerically by using randomly generated pseudo-experiments known as *toy Monte Carlo*. This allows CL_{s+b} and CL_b to be estimated as the fraction of toys for which $t \geq t^{obs}$. Then, if for $\mu = 1$ one obtains $\text{CL}_s \leq \alpha$, it may be stated that the hypothesized new physics phenomenon is excluded with $(1 - \alpha)\text{CL}_s$ confidence level (CL).

The above procedure may be used to set the *observed* exclusion limits. The *expected* limits (and $\pm 1\sigma$ and $\pm 2\sigma$ uncertainty bands) may be obtained by instead generating a large set of toy Monte Carlo and using each of these as if they were real data in calculating CL_s . This allows one to construct a cumulative probability function that yields the median expected upper limit (50% quantile), the $\pm 1\sigma$ bands (16% and 84% quantiles) and the $\pm 2\sigma$ bands (2.5% and 97.5% quantiles).

It should be noted, however, that generating large sets of toy Monte Carlo may become problematic due to the required computational resources. Instead, an asymptotic formulation may be employed [243]. In Ref. [243] Wilk's theorem is used to describe the asymptotic behavior of f with well-defined formulas for both the signal+background and background-only hypotheses. These in turn allow one to express the upper limit as:

$$\text{CL}_s = \alpha = \frac{1 - \Phi(\sqrt{t_\mu})}{\Phi(\sqrt{t_{\mu,A}}) - \sqrt{t_\mu}} \quad (6)$$

where Φ is the cumulative distribution of the standard Gaussian and $t_{\mu,A}$ is the test statistic evaluated with the Asimov dataset⁶. Finally, the median expected CL_s upper limit and the uncertainty bands are given by:

$$\mu_{up+N} = \sigma(\Phi^{-1}(1 - \alpha\Phi(N)) + N) \quad (7)$$

where $N=0$ corresponds to the median expected upper limit and $N = \pm 1$ (± 2) corresponds to the 1σ (2σ) uncertainty bands.

⁶In this context the Asimov dataset may be understood as the expected background with nuisance parameters fixed to their nominal values.

Appendix B

Data-driven Estimation of Backgrounds with Non-prompt or Fake Leptons

The search described in Chapter 5 relies on data-driven techniques for the estimation of the non-prompt or fake lepton background, due to the generally poor modeling of this background with MC simulation. This appendix provides more details on the fake rate measurements, transfer factors and the MC scaling procedures in the semi-data-driven variations of the tight-to-loose method.

B.1 Fake Rate Measurements

Prompt leptons with $\Delta R(\ell\ell) > 0.3$

In the prompt analysis categories, a lepton are treated as independent if there is no other final state lepton in its isolation cone. The fake rate for independent leptons is measured in a QCD-enriched measurement region. Events in the MR are required to have exactly one lepton passing the loose object criteria and at least one jet with $p_T > 50$ GeV that satisfies $\Delta R(\ell, \text{jet}) > 0.7$. For the fake rate measurement of muons, the events are selected at the HLT with single muon paths that were prescaled during data-taking (`HLT_Mu3_PFJet40` and `HLT_Mu8`), while for the electron channel a mixture of pre-scaled single jet triggers were used. The measurements are performed both with MC simulated QCD events (for the closure test) and with observed data events (for the actual background estimate). Since the flavor composition in the MR might differ from that of the AR and SR, which could lead to non-closure of the method, the object requirements are tuned to minimize the flavor dependency. In particular, the threshold of DeepJet tagger score used to veto leptons compatible with B decays (see Tables 5.2 and 5.3) is chosen such as to minimize the flavor dependency without a major cost of signal efficiency.

In the measurement with observed data events, the contamination from electroweak production (leptons from eg. Z or W boson production in association with jets) should be carefully considered in order to avoid biasing the fake rate. Two independent methods are therefore employed to measure the fake rate, both relying on the distribution of the transverse mass, defined as:

$$m_T^{fix}(\ell, p_T^{\text{miss}}) = \sqrt{2 \cdot 35 \cdot p_T^{\text{miss}} (1 - \cos \Delta\phi)} \quad (8)$$

This variable has great discrimination power for QCD and electroweak processes. In this definition, the lepton p_T is fixed to 35 GeV in order to minimize the bias. The first method consists of *data unfolding*; Regions S and L are defined as $0 < m_T^{fix} < 20$ GeV and $70 < m_T^{fix} < 120$ GeV, that are rich in QCD and electroweak processes respectively. The fake rate from the QCD contribution is then measured as:

$$f^{QCD} = \frac{f_S - f_L \cdot r_{V+jets}^{SL}}{1 - r_{V+jets}^{SL}} \quad (9)$$

where $f_{S(L)}$ corresponds to the fake rate measured on observed data in region S (L) of m_T^{fix} and r_{V+jets}^{SL} is given by:

$$r_{V+jets}^{SL} = \left(\frac{N_{V+jets}^S}{N_{V+jets}^L} \right) / \left(\frac{N_{Data}^S}{N_{Data}^L} \right) \quad (10)$$

Here $N_{V+jets}^{S(L)}$ and $N_{Data}^{S(L)}$ denote the number of MC simulated events of electroweak processes and the number of observed data events, respectively, in region S (L) of m_T^{fix} . The second method consists of performing a *simultaneous fit* of MC simulated m_T^{fix} templates to data for passing and failing probes. The fake rate is then extracted from the post-fit QCD MC templates. The final fake rate from data is taken to be the average of the fake rate obtained with the two method. The fake rates are parameterized as function of the lepton p_T and η and the results are shown in Figures 23 and 24.

Prompt leptons with $\Delta R(\ell\ell) < 0.3$

Prompt leptons with $\Delta R(\ell\ell) < 0.3$) are treated as correlated. A dedicated fake rate is measured for the pair, treating the pair as a single object. The measurements of these fake rates are performed with a measurement region defined similar to that of the uncorrelated prompt lepton case, but which instead requires exactly two oppositely charged leptons passing the loose object criteria. The leptons must be

closeby ($\Delta R(\ell\ell < 0.3)$) and contributions from J/ψ and ν decays are vetoed with requirements on the invariant mass of the pair. The fake rate in this case is defined with as denominator the number of events with two loose leptons and as numerator the number of events where the leptons *both* also pass the tight object criteria. In contrast to the case of independent leptons, the contamination of electroweak processes in the correlated lepton MR is very small. Therefore, the electroweak subtraction in the fake rate measurement is performed by simply subtracting the expectation from electroweak MC simulation directly from the number of observed data events. The fake rates are parameterized as function of the lepton pair p_T and shown in Figures 25 and 26.

B

Displaced muon pairs

Similar to the above case, also displaced muon pairs are treated as correlated, which is an assumption already imposed by refitting the muon tracks to a common vertex. The dimuon fake rates is measured in a separate MR, which is defined similar to the displaced 2μ SR (Table 5.5, but requiring a loose dimuon and at least one b tagged jet. Requirements on the hard lepton veto, collinearity and $\log_{10}(\Delta xy/\Delta z)$ are removed and the lower bound on the p_T^{miss}/H_T cut is decreased to 0.2 to increase the number of events populating this MR. Due to b jet requirement, the region is rich in muon pairs from B-hadron decays, which constitute the type of fake lepton background in the displaced SR.

As discussed in Section 5.7.4, muon tracks from pileup tend to appear more isolated, since the PF isolation computation assumes provenance from the PV. As such, the fake lepton contribution from pileup tends to have high fake rates - most objects passing the loose requirements also pass the tight requirements - similar to muons from electroweak processes. Since the tight-to-loose method is not well-suited for the estimation of this background contribution, it is instead estimated from MC simulation, hence should accounted for in the fake rate measurement in a similar fashion as the electroweak processes. Therefore, although the contribution from electroweak processes in the displaced MR is found to be negligible, the small contamination from pileup is simply subtracted from data using the MC estimate. The fake rates are parameterized as function of the lepton p_T and vertex displacement in the transverse plane L_{xy} , and are shown in Figure 27.

Data-driven Estimation of Non-prompt or Fake Leptons

B

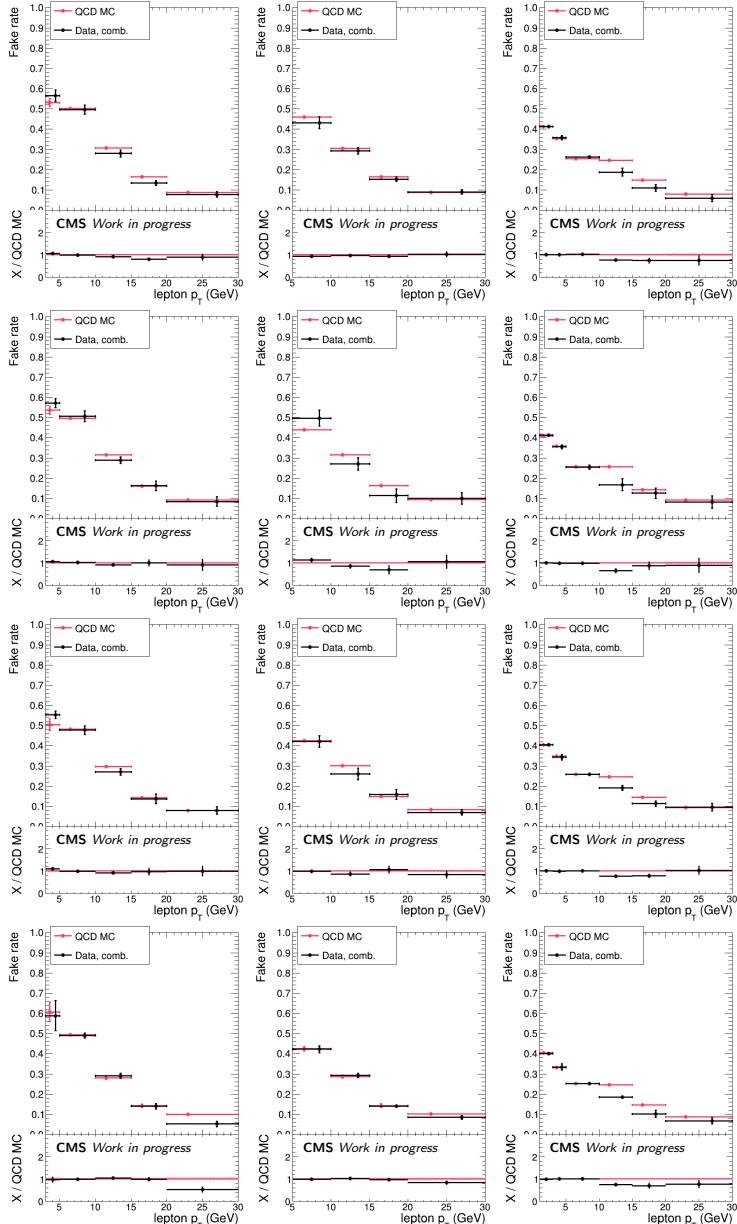


Figure 23: Fake rates measured for prompt independent leptons in the barrel region ($|\eta| < 1.5$). Left column: Muon fake rates based on the HLT_Mu3_PFJet40 trigger, used for $p_T < 10$ GeV. Middle column: Muon fake rates based on the HLT_Mu8 trigger, used for $p_T > 10$ GeV. Right column: Electron fake rates. From top to bottom: 2016APV, 2016, 2017 and 2018 data-taking periods.

B.1 Fake Rate Measurements

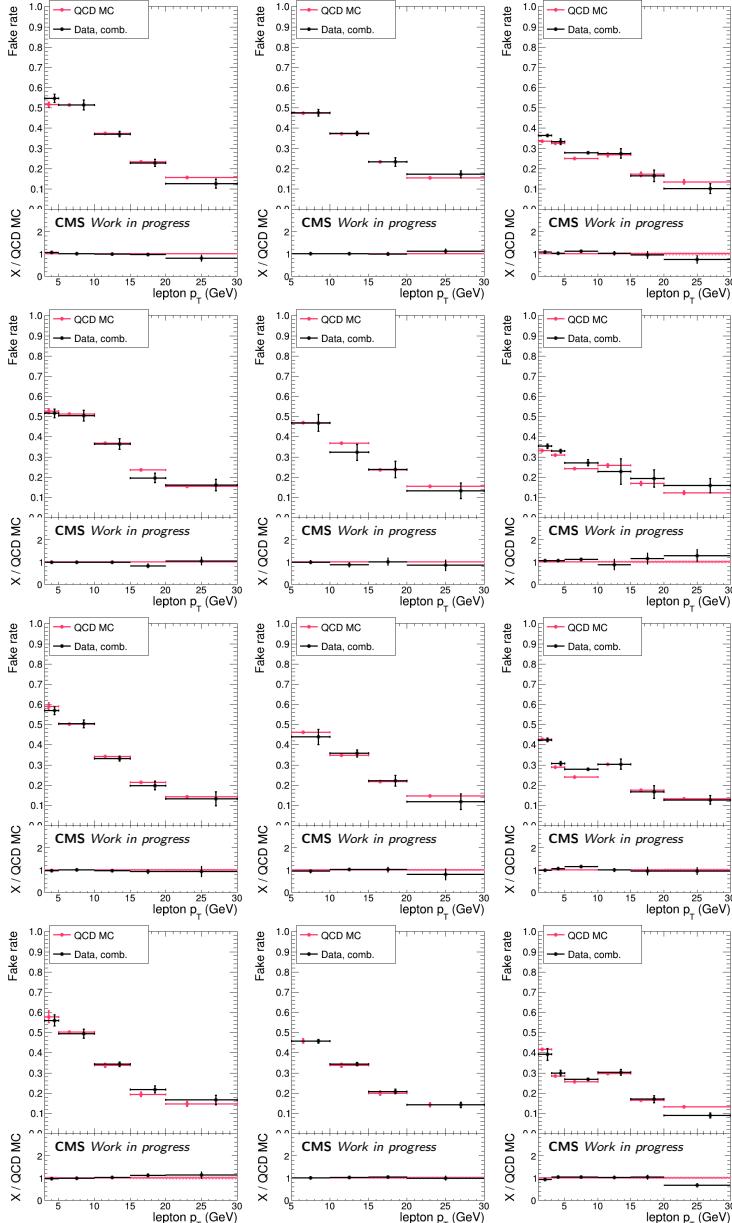


Figure 24: Fake rates measured for prompt independent leptons in the end-cap region ($|\eta| > 1.5$). Left column: Muon fake rates based on the `HLT_Mu3_PFJet40` trigger, used for $p_T < 10$ GeV. Middle column: Muon fake rates based on the `HLT_Mu8` trigger, used for $p_T > 10$ GeV. Right column: Electron fake rates. From top to bottom: 2016APV, 2016, 2017 and 2018 data-taking periods.

Data-driven Estimation of Non-prompt or Fake Leptons

B

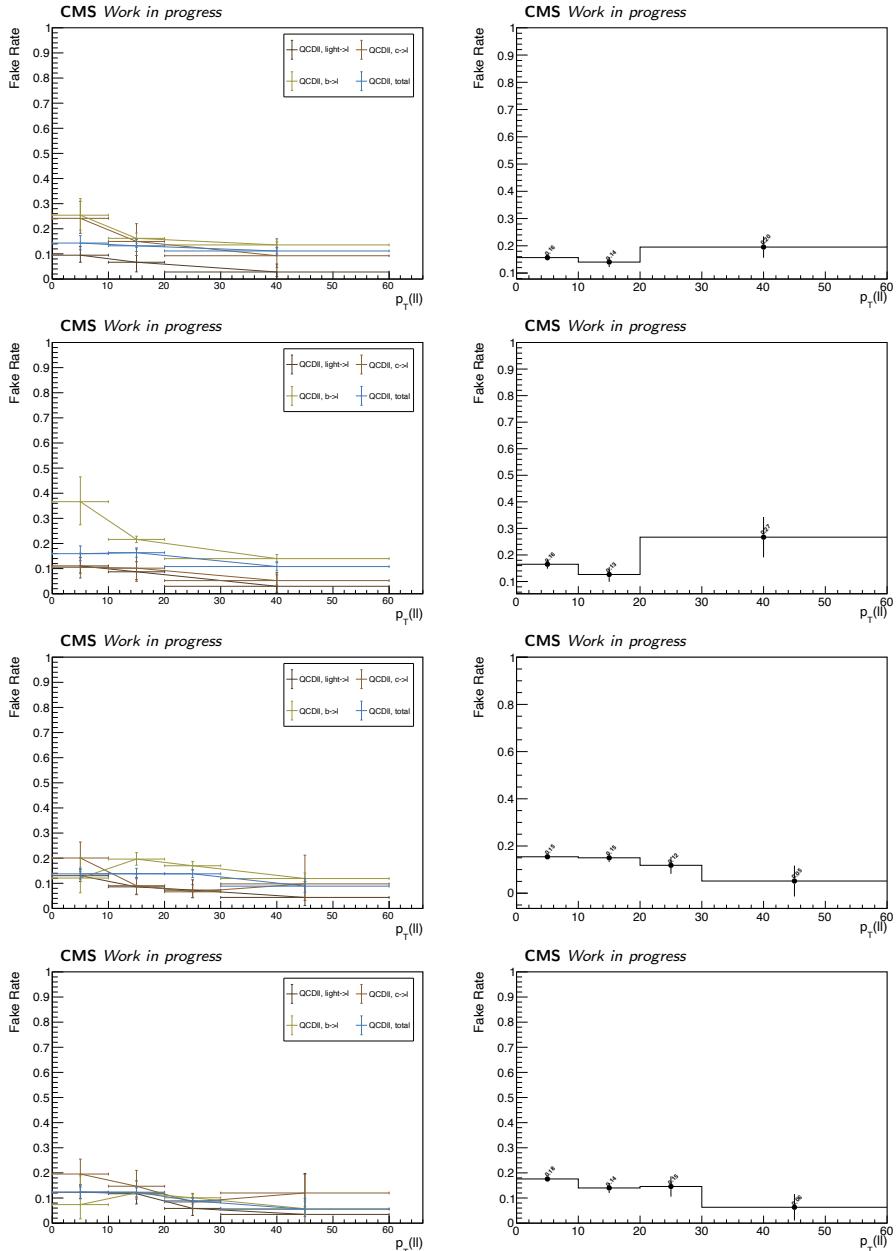


Figure 25: Fake rates for prompt correlated electrons, measured with MC simulation (left) and observed data (right). From top to bottom: 2016APV, 2016, 2017 and 2018 data-taking periods.

B.1 Fake Rate Measurements

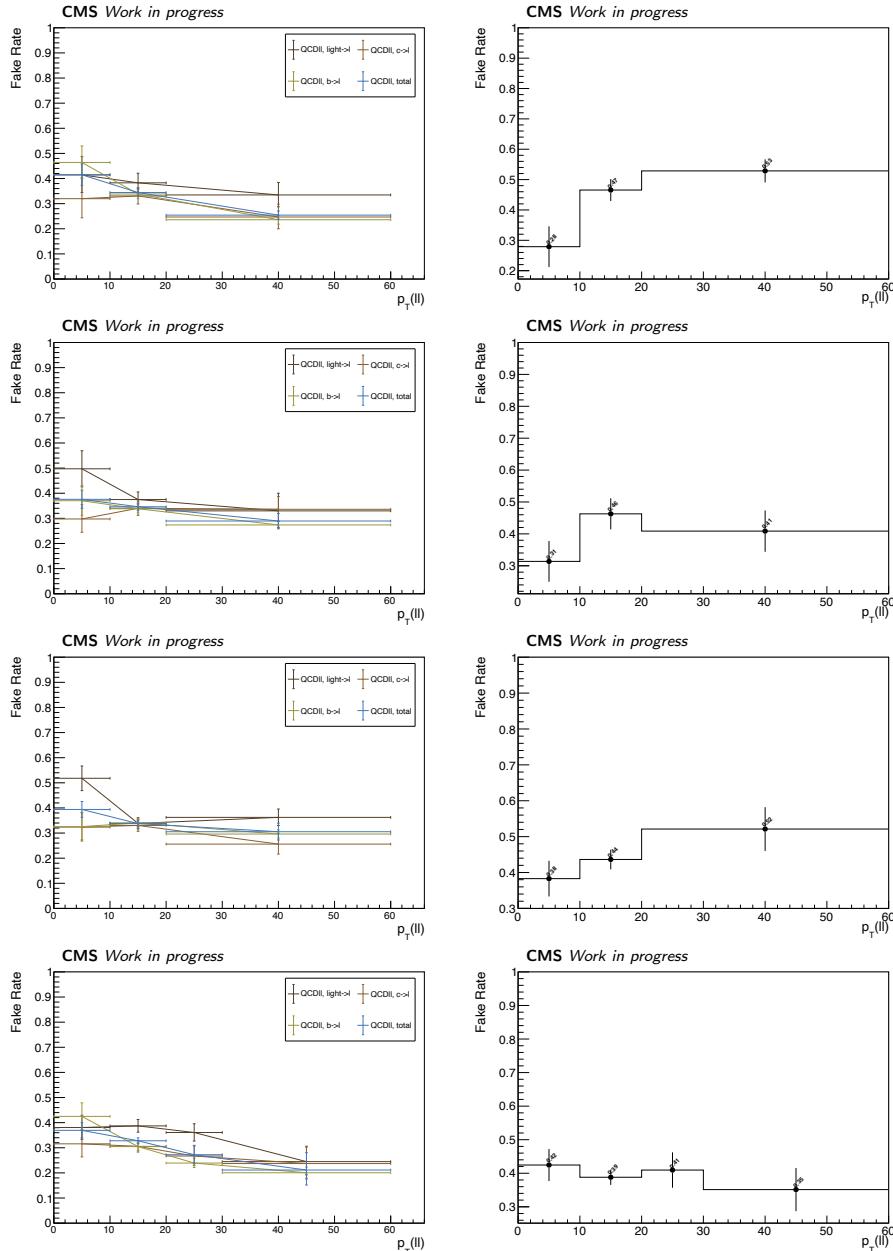


Figure 26: Fake rates for prompt correlated muons, measured with MC simulation (left) and observed data (right). From top to bottom: 2016APV, 2016, 2017 and 2018 data-taking periods.

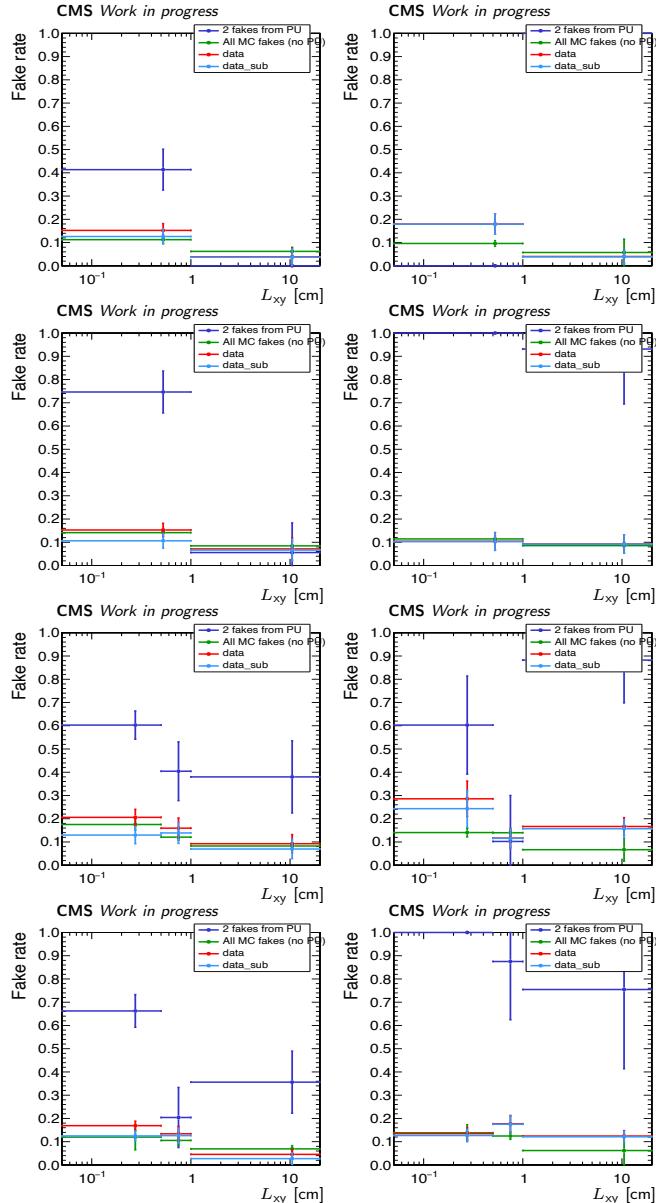


Figure 27: Fake rates for displaced dimuons with $p_T(\ell\ell) < 20$ GeV (left) and $p_T(\ell\ell) > 20$ GeV (right), as function of the transverse displacement of the dimuon vertex with respect to the PV. From top to bottom: 2016APV, 2016, 2017 and 2018 data-taking periods. The contribution from pileup dimuons (dark-blue) is not included in the MC measurement (green), and subtracted from data (red, resulting in light-blue).

B.2 Transfer Factor Derivations

The transfer factors in the tight-to-loose method are used to estimate the fake lepton background contribution in the SR, based on the observed data events in the combined AR and SR. Below follows the derivation for these transfer factors, as well as the simplified expressions corresponding to the semi-data-driven variation of the method, which consists of applying the transfer factors on MC simulated fake lepton events in the AR only.

One-object final state

For events in the prompt 2ℓ category with $\Delta R(\ell\ell) < 0.3$ and events in the displaced 2μ category, the lepton pairs are assumed to be *correlated* and treated as a single object. In both categories the events are separated according to whether the lepton pair that passes the loose object requirements additionally passes (N_T) or fails (N_L) the tight requirements. Both states can receive contributions from events with lepton pairs that are prompt (N_1) or are fake (N_0). N_T and N_L are observed quantities, while N_1 and N_0 are truth-level quantities. By letting ϵ and f be the efficiency for prompt objects (ie. the prompt rate) and the fake rate for fake objects, the yields can be expressed as:

$$\begin{pmatrix} N_T \\ N_L \end{pmatrix} = \begin{pmatrix} \epsilon & f \\ 1 - \epsilon & 1 - f \end{pmatrix} \begin{pmatrix} N_1 \\ N_0 \end{pmatrix} \quad (11)$$

The goal of the method is to estimate the contribution of fake lepton background (N_0) in the SR (N_T), which is hence expressed as $N_T^{bkg} = fN_0$. N_0 can be expressed in terms of the observed quantities by inverting the matrix equation as:

$$\begin{pmatrix} N_1 \\ N_0 \end{pmatrix} = \frac{1}{\epsilon - f} \begin{pmatrix} 1 - f & -f \\ \epsilon - 1 & \epsilon \end{pmatrix} \begin{pmatrix} N_T \\ N_L \end{pmatrix} \quad (12)$$

Substituting the bottom row in the expression for N_T^{bkg} reveals the transfer factors that are to be applied on the observed data events in the SR (W_T) and AR (W_L):

$$W_T = \frac{(\epsilon - 1)f}{\epsilon - f} \quad (13)$$

$$W_L = \frac{\epsilon f}{\epsilon - f} \quad (14)$$

In the semi-data-driven variation, the SR yields are decoupled from the estimation by assuming $\epsilon \gg f$, hence simplifying the transfer factors to:

$$W_T = 0 \quad (15)$$

$$W_L = \frac{f}{1-f} \quad (16)$$

Two-object final state

For events with two *uncorrelated* leptons, or a correlated lepton pair plus an additional independent lepton (a potential scenario in the 3ℓ category), the transfer factors are analogously obtained from a 4×4 matrix equation:

$$\begin{pmatrix} N_{TT} \\ N_{TL} \\ N_{LT} \\ N_{LL} \end{pmatrix} = \begin{pmatrix} \epsilon_1 \epsilon_2 & \epsilon_1 f_2 & f_1 \epsilon_2 & f_1 f_2 \\ \epsilon_1 (1 - \epsilon_2) & \epsilon_1 (1 - f_2) & f_1 (1 - \epsilon_2) & f_1 (1 - f_2) \\ (1 - \epsilon_1) \epsilon_2 & (1 - \epsilon_1) f_2 & (1 - f_1) \epsilon_2 & (1 - f_1) f_2 \\ (1 - \epsilon_1) (1 - \epsilon_2) & (1 - \epsilon_1) (1 - f_2) & (1 - f_1) (1 - \epsilon_2) & (1 - f_1) (1 - f_2) \end{pmatrix} \begin{pmatrix} N_{11} \\ N_{10} \\ N_{01} \\ N_{00} \end{pmatrix} \quad (17)$$

where the first (second) subscript of N by convention refers to the leading (sub-leading) object and ϵ_i (f_i) denotes the prompt (fake) rate of the i -th object. The terms corresponding to the contribution of fake leptons in the SR are therefore:

$$N_{TT}^{bkg} = \epsilon_1 f_2 N_{10} + f_1 \epsilon_2 N_{01} + f_1 f_2 N_{00} \quad (18)$$

Equation 17 is inverted as:

$$\begin{pmatrix} N_{11} \\ N_{10} \\ N_{01} \\ N_{00} \end{pmatrix} = \frac{1}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \begin{pmatrix} (1 - f_1)(1 - f_2) & -(1 - f_1)f_2 & -f_1(1 - f_2) & f_1 f_2 \\ -(1 - f_1)(1 - \epsilon_2) & (1 - f_1)\epsilon_2 & f_1(1 - \epsilon_2) & -f_1 \epsilon_2 \\ -(1 - \epsilon_1)(1 - f_2) & (1 - \epsilon_1)f_2 & \epsilon_1(1 - f_2) & -\epsilon_1 f_2 \\ (1 - \epsilon_1)(1 - \epsilon_2) & -(1 - \epsilon_1)\epsilon_2 & -\epsilon_1(1 - \epsilon_2) & \epsilon_1 \epsilon_2 \end{pmatrix} \begin{pmatrix} N_{pp} \\ N_{pf} \\ N_{fp} \\ N_{ff} \end{pmatrix} \quad (19)$$

Substituting the expressions for N_{10} , N_{01} and N_{00} in Equation 18 and grouping the terms, yields the transfer factors:

$$\begin{aligned} W_{T,T} &= \frac{(1 - \epsilon_1)(1 - \epsilon_2)f_1 f_2}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \\ &\quad - \frac{\epsilon_1 f_2 (1 - f_1)(1 - \epsilon_2)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \\ &\quad - \frac{f_1 \epsilon_2 (1 - \epsilon_1)(1 - f_2)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \end{aligned} \quad (20)$$

(21)

$$W_{T,L} = \frac{\epsilon_1 f_2 \epsilon_2 (1 - f_1)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \quad (22)$$

$$W_{L,T} = \frac{\epsilon_1 f_1 \epsilon_2 (1 - f_2)}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \quad (23)$$

$$W_{L,L} = \frac{-\epsilon_1 \epsilon_2 f_1 f_2}{(\epsilon_1 - f_1)(\epsilon_2 - f_2)} \quad (24)$$

In the semi-data-driven variation (with the assumption $\epsilon \gg f$) these simplify to:

$$W_{T,T} = 0 \quad (25)$$

$$W_{T,L} = \frac{f_2}{1 - f_2} \quad (26)$$

$$W_{L,T} = \frac{f_1}{1 - f_1} \quad (27)$$

$$W_{L,L} = \frac{-f_1 f_2}{(1 - f_2)(1 - f_2)} \quad (28)$$

Three-object final state

The transfer factors for events with three independent leptons are derived in the same fashion, this time based on a 9×9 matrix equation. Although the expressions for the full data-driven method are not shown here due to the lengthy permutations, the transfer factors in the semi-data-driven variation are straightforward; Events with a single lepton failing the tight object requirements receive a weight $f/(1 - f)$, events with two failing leptons receive a weight $-f_i f_j / ((1 - f_i)(1 - f_j))$, and events with all leptons failing receive a weight equal to the product of the $f/(1 - f)$ factors evaluated on the three leptons.

B.3 Semi-data-driven Scale Factors

An overview of the (semi-)data-driven variations used for the fake lepton background estimates in each search category is shown in Table 6. In the semi-data-driven variation of the tight-to-loose method, the MC simulated fake lepton background events in the AR are scaled to the observed data, in order to let the latter govern the shape and/or normalization of the $m_{\ell\ell}$ templates. The scaling procedures are described in the following.

As shown in the un-scaled AR distributions of Figure 5.28, the MC templates in the higher-MET prompt 2ℓ regions poorly describe the observed data. Therefore, in these regions the fake lepton MC estimate is scaled to the observed data

Category	low-MET	medium-MET	high-MET	ultra-MET
prompt $2e$	–	semi-dd <i>bin-by-bin</i>	semi-dd <i>bin-by-bin</i>	semi-dd <i>bin-by-bin</i>
prompt 2μ	dd	semi-dd <i>bin-by-bin</i>	semi-dd <i>bin-by-bin</i>	semi-dd <i>bin-by-bin</i>
prompt 3ℓ	semi-dd <i>inclusive</i>		semi-dd <i>inclusive</i>	
displaced 2μ	dd		dd	

Table 6: Summary of the non-prompt lepton background estimations used for each of the analysis regions. In the fully data-driven (dd) case, transfer factors are applied on observed data events in the combined AR and SR. In the semi-dd case, transfer factors are applied on AR fake lepton MC events, scaled to data either in each $m_{\ell\ell}$ bin independently (bin-by-bin) or inclusive in $m_{\ell\ell}$.

B

independently in each bin of the $m_{\ell\ell}$ distribution, accounting also for the prompt lepton contribution in the AR:

$$\text{Scale Factor} = \frac{\text{AR data} - \text{AR prompt lepton MC}}{\text{AR fake lepton MC}} \quad (29)$$

For robustness against statistical fluctuations, the scaling is performed in the inclusive higher-MET region, ie. for $p_T^{\text{miss,corr}} > 200$ GeV. To obtain the fake lepton MC distribution normalised in each of the individual MET bins, a rate factor is applied to the scaled estimates in the inclusive higher-MET region, defined as:

$$\text{Rate Factor} = \frac{\text{AR fake lepton MC in reference MET bin}}{\text{AR fake lepton MC in inclusive MET bin}} \quad (30)$$

where the numerator and denominator are sums over the full range of $m_{\ell\ell}$ in the given MET bin. It should be noted that this rate factor only applies to prompt 2ℓ region, since the higher-MET 3ℓ region consists of only a single MET bin. The measured scale factors and rate factors for the 2ℓ regions are shown in Tables 7 and 8.

For the prompt 3ℓ AR region, the semi-data-driven SF is measured inclusively in $m_{\ell\ell}$, since the templates from fake lepton MC already well describe the observed data (Figure 5.28). However, it was found that the semi-data-driven estimate could still lead to negative predictions in the SR, since the fraction of events with one, two and three leptons failing the tight object criteria is poorly estimated with MC simulation. Therefore, the SF is measured separately in side-bands of the number of leptons failing the tight object requirements. By scaling the non-prompt

B.3 Semi-data-driven Scale Factors

MC independently in each side-band, the fraction of negatively weighted events (with two loose-not-tight leptons) is corrected with the observed data, guaranteeing positive-definite SR predictions in all bins of $m_{\ell\ell}$. The scale factors for the 3ℓ regions are shown in Table 9. The AR distributions after scaling the fake lepton MC templates are shown in Figure 28.

$m_{\ell\ell}$ bin	2016APV	2016	2017	2018
$0.1 < m_{\ell\ell} < 4$ GeV	7.291 (1.171)	4.376 (1.013)	0.750 (1.150)	1.332 (1.942)
$4 < m_{\ell\ell} < 10$ GeV	1.091 (0.866)	1.734 (0.555)	0.605 (1.378)	4.515 (1.386)
$10 < m_{\ell\ell} < 20$ GeV	0.496 (1.758)	0.789 (0.677)	0.340 (1.290)	0.848 (0.931)
$20 < m_{\ell\ell} < 30$ GeV	0.266 (0.351)	1.000 (1.438)	0.551 (1.703)	1.188 (1.162)
$30 < m_{\ell\ell} < 50$ GeV	1.000 (0.436)	0.833 (1.472)	0.422 (0.850)	1.304 (1.223)

Table 7: The semi-data-driven SFs derived from the prompt 2ℓ -AR for muons (electrons), with the higher MET bins merged.

MET bin	2016APV	2016	2017	2018
medium-MET	0.454 (0.486)	0.318 (0.456)	0.552 (0.454)	0.541 (0.477)
high-MET	0.314 (0.248)	0.393 (0.346)	0.195 (0.286)	0.266 (0.256)
ultra-MET	0.232 (0.266)	0.289 (0.198)	0.252 (0.261)	0.193 (0.267)

Table 8: The semidata-driven rate factors derived from the prompt 2ℓ -AR for muons (electrons).

AR side-band	2016APV		2016		2017		2018	
	low-MET	high-MET	low-MET	high-MET	low-MET	high-MET	low-MET	high-MET
1 LNT	1.000	1.331	0.668	1.221	1.436	0.667	1.189	1.276
2 LNT	0.276	1.338	1.133	1.088	1.022	1.382	0.421	1.112
3 LNT	1.000	2.997	1.000	1.403	1.000	2.263	1.000	2.343

Table 9: The semi-data-driven SFs derived from the prompt 3ℓ -ARs. In contrast to the prompt 2ℓ regions, the SFs in the 3ℓ -ARs are measured inclusively in $m_{\ell\ell\ell}$, but separately in side-bands of the number of loose-not-tight (LNT) leptons. For the low-MET 3 LNT side-band no data events were observed, hence the semi-data-driven SF is fixed to 1.

Data-driven Estimation of Non-prompt or Fake Leptons

B

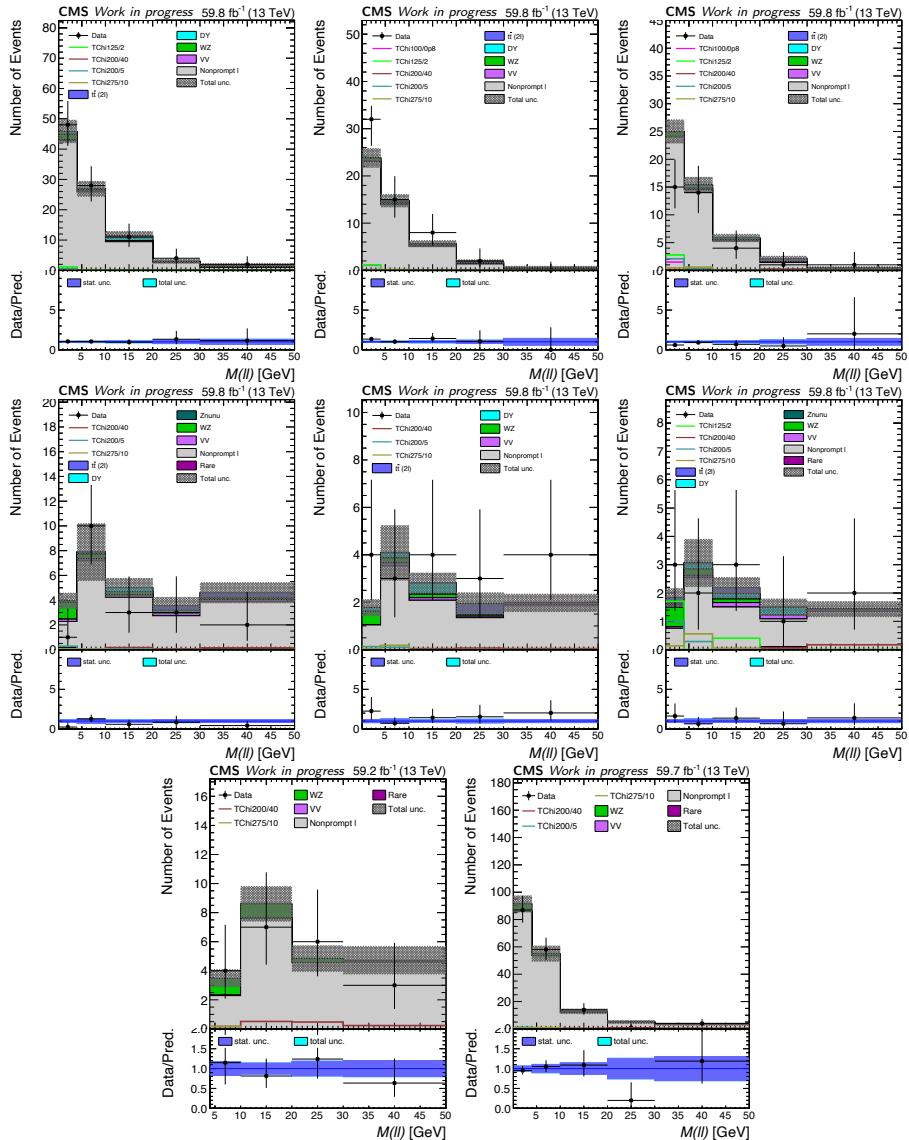


Figure 28: The $m_{\ell\ell}$ templates in the prompt 2ℓ -AR for electrons (top row) and muons (middle row), and in the prompt 3ℓ -AR (bottom row), corresponding to the 2018 data-taking period. From left to right the plots correspond to the medium, high and ultra-MET bins for the 2ℓ categories and the low and high-MET bins for the 3ℓ category. The fake lepton background estimates from MC are scaled according to the semi-data-driven variation of the tight-to-loose method. The plots can be compared to the un-scaled versions in Figure 5.28.

Appendix C

Features for Multi-variate Electron Identification in the CMS Phase-2 Level-1 Trigger

This appendix provides a description of all cluster, track and matching features that were considered for the multi-variate electron identification in the Phase 2 Level-1 Trigger. Different sets of features have been used in different algorithms:

- The 5 (cluster) features labeled with \square are used for the PU identification algorithm (Section 7.2.2), which is evaluated for all 3D clusters constructed by the HGCAL-TPG.
- The 9 (cluster) features labeled with \diamond are used for the stand-alone e/γ identification (Section 7.2.2), part of the TkElectron algorithm.
- The 27 features labeled with \star are used for the first version of the Composite ID (Section 7.3.2), based on simulation.
- The 18 features labeled with \heartsuit will be available in the Correlator Trigger firmware and can be used for algorithms run in the Correlator Trigger boards. The other features are not available, because they cannot be implemented in the backends that generate the TPs (cluster and tracks) due to limited latency and FPGA resources.
- The 11 features labeled with \triangle are chosen for the final Composite ID that was implemented in firmware (Section 7.3.4).

Cluster Features (1/2)

Feature Name	Description	Label
FIRSTLAYER	First HGCAL layer of the cluster	◊★⚙
LASTLAYER	Last HGCAL layer of the cluster	⚙
MAXLAYER	HGCAL layer with maximum energy deposit	◊★
SHOWERLENGTH	LASTLAYER - FIRSTLAYER + 1	◊★⚙
CORESHOWERLENGTH	Maximum number of consecutive HGCAL layers	◊★⚙
SZZ	Energy-weighted RMS of the z-coordinate of the trigger cells of the cluster	□◊★⚙
SRRTOT	Energy-weighted RMS of the r-coordinate of the trigger cells of the cluster	◊★⚙△
SEETOT	Energy-weighted RMS of the η -coordinate of the trigger cells of the cluster	◊★⚙
SPPTOT	Energy-weighted RMS of the ϕ -coordinate of the trigger cells of the cluster	□◊★⚙
SRRMEAN	Energy weighted average of layers computed for each layer for trigger cells within 5 cm of max energy trigger	◊★

Cluster Features (2/2)

Feature Name	Description	Label
EMAXE	Energy of layer with maximum energy	□ ★
MEANZ	Center of the cluster in the z-coordinate	★ ⚙ ▲
LAYER10	Number of layers at which 10% of the cluster energy is deposited	★
LAYER50	Number of layers at which 50% of the cluster energy is deposited	□ ★
LAYER90	Number of layers at which 90% of the cluster energy is deposited	★
NTC67	Number of trigger cells at which 67% of the cluster energy is deposited	□ ★
NTC90	Number of trigger cells at which 90% of the cluster energy is deposited	★
HOE	Cluster energy in HCAL divided by cluster energy in ECAL	★ ⚙ ▲
SEEMAX	Maximum energy-weighted RMS of the η -coordinate of the trigger cells of all layers	★
SPPMAX	Maximum energy-weighted RMS of the ϕ -coordinate of the trigger cells of all layers	★
SRRMAX	Maximum energy-weighted RMS of the r-coordinate of the trigger cells of all layers	★

Track Features

Feature Name	Description	Label
TKCHI2	χ^2 of the overall track fit	★
TKCHI2RPHI	χ^2 of the track fit in the $r\phi$ plane	★△
TKCHI2RZ	χ^2 of the track fit in the rz plane	★△
TKCHI2BEND	χ^2 between the track fit p_T and the coarse p_T measurement from the bend of each stub	★△
TKZ0	z-coordinate of the track at vertex	★★
TKNSTUBS	Number of stubs of the track	★★△
TKPT	p_T of the track	★★△

C

Matching Features

Feature Name	Description	Label
DETA	Difference in η between the track and cluster at the calorimeter surface	★★△
DPHI	Difference in ϕ between the track and cluster at the calorimeter surface	★★△
DPT	Ratio of track p_T and cluster p_T	★★△

Bibliography

- [1] S. L. Glashow, “Partial-symmetries of weak interactions”, *Nuclear Physics* **22** (1961) 579, [doi:10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2).
- [2] S. Weinberg, “A Model of Leptons”, *Phys. Rev. Lett.* **19** (1967) 1264, [doi:10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).
- [3] S. L. Glashow, J. Iliopoulos, and L. Maiani, “Weak Interactions with Lepton-Hadron Symmetry”, *Phys. Rev. D* **2** (1970) 1285, [doi:10.1103/PhysRevD.2.1285](https://doi.org/10.1103/PhysRevD.2.1285).
- [4] M. Gell-Mann, “A schematic model of baryons and mesons”, *Phys. Lett.* **8** (1964) 214, [doi:10.1016/S0031-9163\(64\)92001-3](https://doi.org/10.1016/S0031-9163(64)92001-3).
- [5] G. Zweig, “An SU(3) model for strong interaction symmetry and its breaking. Version 2”, [doi:10.17181/CERN-TH-412](https://doi.org/10.17181/CERN-TH-412).
- [6] 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).
- [7] 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).
- [8] J. Wess and B. Zumino, “Supergauge transformations in four dimensions”, *Nucl. Phys. B* **70** (1974) 39, [doi:10.1016/0550-3213\(74\)90355-1](https://doi.org/10.1016/0550-3213(74)90355-1).
- [9] H. Nilles, “Supersymmetry, supergravity and particle physics”, *Physics Reports* **110** (1984) 1, [doi:10.1016/0370-1573\(84\)90008-5](https://doi.org/10.1016/0370-1573(84)90008-5).

Bibliography

- [10] H. E. Haber and G. L. Kane, “The search for supersymmetry: Probing physics beyond the standard model”, *Phys. Rept.* **117** (1985) 75, [doi:10.1016/0370-1573\(85\)90051-1](https://doi.org/10.1016/0370-1573(85)90051-1).
- [11] R. Barbieri, S. Ferrara, and C. A. Savoy, “Gauge models with spontaneously broken local supersymmetry”, *Phys. Lett. B* **119** (1982) 343, [doi:10.1016/0370-2693\(82\)90685-2](https://doi.org/10.1016/0370-2693(82)90685-2).
- [12] S. Dawson, E. Eichten, and C. Quigg, “Search for supersymmetric particles in hadron-hadron collisions”, *Phys. Rev. D* **31** (1985) 1581, [doi:10.1103/PhysRevD.31.1581](https://doi.org/10.1103/PhysRevD.31.1581).
- [13] P. Langacker, “The Standard Model and Beyond”. Taylor & Francis, first edition, 2010. ISBN 978-1-4200-7906-7.
- [14] M. Thomson, “Modern particle physics”. Cambridge University Press, New York, first edition, 2013. ISBN 978-1-107-03426-6.
- [15] D. Griffiths, “Introduction to elementary particles”. Wiley, second edition, 2008. ISBN 978-3-527-40601-2.
- [16] ParticleDataGroup Collaboration, “Review of Particle Physics: 2022”, *PTEP* **2022** (2022) 083C01, [doi:10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [17] Y. Fukuda et al., “Evidence for oscillation of atmospheric neutrinos”, *Phys. Rev. Lett.* **81** (1998) 1562, [doi:10.1103/physrevlett.81.1562](https://doi.org/10.1103/physrevlett.81.1562).
- [18] A. Deur, S. J. Brodsky, and G. F. de Téramond, “The QCD running coupling”, *Progress in Particle and Nuclear Physics* **90** (2016) 1, [doi:10.1016/j.ppnp.2016.04.003](https://doi.org/10.1016/j.ppnp.2016.04.003).
- [19] S. P. Martin, “A supersymmetry primer”, *ASDHE* (1998) 1, [doi:10.1142/9789812839657_0001](https://doi.org/10.1142/9789812839657_0001), [arXiv:hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356).
- [20] D. J. Gross and F. Wilczek, “Ultraviolet behavior of non-abelian gauge theories”, *Phys. Rev. Lett.* **30** (1973) 1343, [doi:10.1103/PhysRevLett.30.1343](https://doi.org/10.1103/PhysRevLett.30.1343).
- [21] H. D. Politzer, “Reliable perturbative results for strong interactions?”, *Phys. Rev. Lett.* **30** (1973) 1346, [doi:10.1103/PhysRevLett.30.1346](https://doi.org/10.1103/PhysRevLett.30.1346).
- [22] J. C. Collins and D. E. Soper, “The theorems of perturbative QCD”, *Annual Review of Nuclear and Particle Science* **37** (1987) 383, [doi:10.1146/annurev.ns.37.120187.002123](https://doi.org/10.1146/annurev.ns.37.120187.002123).
- [23] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, “Parton distributions for the LHC”, *Eur. Phys. J. C* **63** (2009) 189–285, [doi:10.1140/epjc/s10052-009-1072-5](https://doi.org/10.1140/epjc/s10052-009-1072-5), [arXiv:0901.0002](https://arxiv.org/abs/0901.0002).

- [24] CMS Collaboration, “CMS Standard Model summary plots”. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsCombined>.
- [25] V. C. Rubin, J. Ford, W. K., and N. Thonnard, “Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 ($R=4\text{kpc}$) to UGC 2885 ($R=122\text{kpc}$).”, *The Astrophysical Journal* **238** (1980) 471, [doi:10.1086/158003](https://doi.org/10.1086/158003).
- [26] K. G. Begeman, A. H. Broeils, and R. H. Sanders, “Extended rotation curves of spiral galaxies: Dark haloes and modified dynamics”, *Mon. Not. Roy. Astron. Soc.* **249** (1991) 523, [doi:10.1093/mnras/249.3.523](https://doi.org/10.1093/mnras/249.3.523).
- [27] A. D. Sakharov, “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe”, *Pisma Zh. Eksp. Teor. Fiz.* **5** (1967) 32–35, [doi:10.1070/PU1991v034n05ABEH002497](https://doi.org/10.1070/PU1991v034n05ABEH002497).
- [28] W. Bernreuther, “Cp violation and baryogenesis”, 2002.
- [29] M. Dine and A. Kusenko, “Origin of the matter-antimatter asymmetry”, *Rev. Mod. Phys.* **76** (2003) 1, [doi:10.1103/revmodphys.76.1](https://doi.org/10.1103/revmodphys.76.1).
- [30] G. 't Hooft, “Computation of the quantum effects due to a four-dimensional pseudoparticle”, *Phys. Rev. D* **14** (1976) 3432, [doi:10.1103/PhysRevD.14.3432](https://doi.org/10.1103/PhysRevD.14.3432).
- [31] J. M. Pendlebury et al., “Revised experimental upper limit on the electric dipole moment of the neutron”, *Phys. Rev. D* **92** (2015) 092003, [doi:10.1103/PhysRevD.92.092003](https://doi.org/10.1103/PhysRevD.92.092003).
- [32] J. E. Kim and G. Carosi, “Axions and the Strong CP Problem”, *Rev. Mod. Phys.* **82** (2010) 557–602, [doi:10.1103/RevModPhys.82.557](https://doi.org/10.1103/RevModPhys.82.557), [arXiv:0807.3125](https://arxiv.org/abs/0807.3125). [Erratum: Rev.Mod.Phys. 91, 049902 (2019)].
- [33] D. Chung et al., “The soft supersymmetry-breaking lagrangian: theory and applications”, *Physics Reports* **407** (2005) 1, [doi:10.1016/j.physrep.2004.08.032](https://doi.org/10.1016/j.physrep.2004.08.032).
- [34] D. Eliezer and R. Woodard, “The problem of nonlocality in string theory”, *Nucl. Phys. B* **325** (1989) 389, [doi:10.1016/0550-3213\(89\)90461-6](https://doi.org/10.1016/0550-3213(89)90461-6).
- [35] R. Barbieri and G. Giudice, “Upper bounds on supersymmetric particle masses”, *Nucl. Phys. B* **306** (1988) 63, [doi:10.1016/0550-3213\(88\)90171-X](https://doi.org/10.1016/0550-3213(88)90171-X).
- [36] S. Dimopoulos and H. Georgi, “Softly broken supersymmetry and $SU(5)$ ”, *Nucl. Phys. B* **193** (1981) 150, [doi:10.1016/0550-3213\(81\)90522-8](https://doi.org/10.1016/0550-3213(81)90522-8).

Bibliography

- [37] E. Witten, “Dynamical breaking of supersymmetry”, *Nucl. Phys. B* **188** (1981) 513, [doi:10.1016/0550-3213\(81\)90006-7](https://doi.org/10.1016/0550-3213(81)90006-7).
- [38] M. Dine, W. Fischler, and M. Srednicki, “Supersymmetric technicolor”, *Nucl. Phys. B* **189** (1981) 575, [doi:10.1016/0550-3213\(81\)90582-4](https://doi.org/10.1016/0550-3213(81)90582-4).
- [39] S. Dimopoulos and S. Raby, “Supercolor”, *Nucl. Phys. B* **192** (1981) 353, [doi:10.1016/0550-3213\(81\)90430-2](https://doi.org/10.1016/0550-3213(81)90430-2).
- [40] N. Sakai, “Naturalness in supersymmetric GUTS”, *Z. Phys. C* **11** (1981) 153, [doi:10.1007/BF01573998](https://doi.org/10.1007/BF01573998).
- [41] R. K. Kaul and P. Majumdar, “Cancellation of quadratically divergent mass corrections in globally supersymmetric spontaneously broken gauge theories”, *Nucl. Phys. B* **199** (1982) 36, [doi:10.1016/0550-3213\(82\)90565-X](https://doi.org/10.1016/0550-3213(82)90565-X).
- [42] B. de Carlos and J. A. Casas, “One loop analysis of the electroweak breaking in supersymmetric models and the fine tuning problem”, *Phys. Lett. B* **309** (1993) 320–328, [doi:10.1016/0370-2693\(93\)90940-J](https://doi.org/10.1016/0370-2693(93)90940-J), [arXiv:hep-ph/9303291](https://arxiv.org/abs/hep-ph/9303291).
- [43] R. Barbieri and D. Pappadopulo, “S-particles at their naturalness limits”, *JHEP* **10** (2009) 061, [doi:10.1088/1126-6708/2009/10/061](https://doi.org/10.1088/1126-6708/2009/10/061), [arXiv:0906.4546](https://arxiv.org/abs/0906.4546).
- [44] D. Hooper, “Tasi 2008 lectures on dark matter”, [arXiv:0901.4090](https://arxiv.org/abs/0901.4090).
- [45] N. Polonsky, “Supersymmetry: Structure and phenomena. Extensions of the standard model”, volume 68. 2001. [doi:10.1007/3-540-44642-7](https://doi.org/10.1007/3-540-44642-7).
- [46] Muon $g - 2$ Collaboration, “Measurement of the positive muon anomalous magnetic moment to 0.46 ppm”, *Phys. Rev. Lett.* **126** (2021) 141801, [doi:10.1103/PhysRevLett.126.141801](https://doi.org/10.1103/PhysRevLett.126.141801), [arXiv:2104.03281](https://arxiv.org/abs/2104.03281).
- [47] Muon $g - 2$ Collaboration, “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm”, *Phys. Rev. Lett.* **131** (2023) 161802, [doi:10.1103/PhysRevLett.131.161802](https://doi.org/10.1103/PhysRevLett.131.161802), [arXiv:2308.06230](https://arxiv.org/abs/2308.06230).
- [48] S. Borsanyi et al., “Leading hadronic contribution to the muon magnetic moment from lattice QCD”, *Nature* **593** (2021) 51, [doi:10.1038/s41586-021-03418-1](https://doi.org/10.1038/s41586-021-03418-1).
- [49] G. Colangelo et al., “Prospects for precise predictions of a_μ in the standard model”, [arXiv:2203.15810](https://arxiv.org/abs/2203.15810).
- [50] M. Endo, K. Hamaguchi, S. Iwamoto, and T. Kitahara, “Supersymmetric interpretation of the muon $g - 2$ anomaly”, *JHEP* **2021** (2021)

[doi:10.1007/jhep07\(2021\)075](https://doi.org/10.1007/jhep07(2021)075).

- [51] CDF Collaboration, “High-precision measurement of the W boson mass with the CDF II detector”, *Science* **376** (2022) 170–176, [doi:10.1126/science.abk1781](https://doi.org/10.1126/science.abk1781).
- [52] S. Heinemeyer, “CDF measurement of m_w : Theory implications”, [arXiv:2207.14809](https://arxiv.org/abs/2207.14809).
- [53] M. Maniatis, “The next-to-minimal supersymmetric extension of the standard model reviewed”, [doi:10.1142/s0217751x10049827](https://doi.org/10.1142/s0217751x10049827).
- [54] A. Kaminska, G. G. Ross, and K. Schmidt-Hoberg, “Non-universal gaugino masses and fine tuning implications for SUSY searches in the MSSM and the GNMSSM”, *JHEP* **2013** (2013) [doi:10.1007/jhep11\(2013\)209](https://doi.org/10.1007/jhep11(2013)209).
- [55] S. Deser and B. Zumino, “Broken supersymmetry and supergravity”, *Phys. Rev. Lett.* **38** (1977) 1433, [doi:10.1103/PhysRevLett.38.1433](https://doi.org/10.1103/PhysRevLett.38.1433).
- [56] G. Giudice and R. Rattazzi, “Theories with gauge-mediated supersymmetry breaking”, *Physics Reports* **322** (1999) 419, [doi:10.1016/S0370-1573\(99\)00042-3](https://doi.org/10.1016/S0370-1573(99)00042-3).
- [57] N. Arkani-Hamed and S. Dimopoulos, “Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC”, *JHEP* **2005** (2005) 073, [doi:10.1088/1126-6708/2005/06/073](https://doi.org/10.1088/1126-6708/2005/06/073).
- [58] N. Arkani-Hamed, S. Dimopoulos, G. Giudice, and A. Romanino, “Aspects of split supersymmetry”, *Nucl. Phys. B* **709** (2005) 3, [doi:10.1016/j.nuclphysb.2004.12.026](https://doi.org/10.1016/j.nuclphysb.2004.12.026).
- [59] G. Giudice and A. Romanino, “Erratum to: “split supersymmetry” [nucl. phys. b 699 (2004) 65]”, *Nucl. Phys. B* **706** (2005) 487, [doi:10.1016/j.nuclphysb.2004.11.048](https://doi.org/10.1016/j.nuclphysb.2004.11.048).
- [60] N. Arkani-Hamed et al., “Simply unnatural supersymmetry”, 2012.
- [61] A. Arvanitaki, N. Craig, S. Dimopoulos, and G. Villadoro, “Mini-split”, *JHEP* **2013** (2013) [doi:10.1007/jhep02\(2013\)126](https://doi.org/10.1007/jhep02(2013)126).
- [62] G. L. Kane, C. Kolda, L. Roszkowski, and J. D. Wells, “Study of constrained minimal supersymmetry”, *Phys. Rev. D* **49** (1994) 6173, [doi:10.1103/physrevd.49.6173](https://doi.org/10.1103/physrevd.49.6173).
- [63] H. Okawa, “Interpretations of SUSY searches in ATLAS with simplified models”, [arXiv:1110.0282](https://arxiv.org/abs/1110.0282).
- [64] S. Chatrchyan et al., “Interpretation of searches for supersymmetry with simplified models”, *Phys. Rev. D* **88** (2013)

Bibliography

[doi:10.1103/physrevd.88.052017](https://doi.org/10.1103/physrevd.88.052017).

[65] LHC SUSY Cross Section Working Group Collaboration, “SUSY Cross Sections”. <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections>.

[66] C. F. Berger, J. S. Gainer, J. L. Hewett, and T. G. Rizzo, “Supersymmetry without prejudice”, *JHEP* **2009** (2009) 023, [doi:10.1088/1126-6708/2009/02/023](https://doi.org/10.1088/1126-6708/2009/02/023).

[67] V. Khachatryan et al., “Phenomenological MSSM interpretation of CMS searches in pp collisions at $\sqrt{s} = 7$ and 8 TeV”, *JHEP* **2016** (2016) [doi:10.1007/jhep10\(2016\)129](https://doi.org/10.1007/jhep10(2016)129).

[68] G. Aad et al., “Summary of the ATLAS experiment’s sensitivity to supersymmetry after LHC run 1 — interpreted in the phenomenological MSSM”, *JHEP* **2015** (2015) [doi:10.1007/jhep10\(2015\)134](https://doi.org/10.1007/jhep10(2015)134).

[69] G. D’Ambrosio, G. Giudice, G. Isidori, and A. Strumia, “Minimal flavour violation: an effective field theory approach”, *Nucl. Phys. B* **645** (2002) 155, [doi:10.1016/S0550-3213\(02\)00836-2](https://doi.org/10.1016/S0550-3213(02)00836-2).

[70] J. M. Butterworth, A. R. Davison, M. Rubin, and G. P. Salam, “Jet substructure as a new higgs-search channel at the large hadron collider”, *Phys. Rev. Lett.* **100** (2008) [doi:10.1103/physrevlett.100.242001](https://doi.org/10.1103/physrevlett.100.242001).

[71] L. Randall and D. Tucker-Smith, “Dijet searches for supersymmetry at the large hadron collider”, *Phys. Rev. Lett.* **101** (2008) [doi:10.1103/physrevlett.101.221803](https://doi.org/10.1103/physrevlett.101.221803).

[72] C. Lester and D. Summers, “Measuring masses of semi-invisibly decaying particle pairs produced at hadron colliders”, *Phys. Lett. B* **463** (1999) 99, [doi:10.1016/s0370-2693\(99\)00945-4](https://doi.org/10.1016/s0370-2693(99)00945-4).

[73] D. R. Tovey, “On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders”, *JHEP* **2008** (2008) 034, [doi:10.1088/1126-6708/2008/04/034](https://doi.org/10.1088/1126-6708/2008/04/034).

[74] CMS Collaboration, “Inclusive search for squarks and gluinos in pp collisions at $\sqrt{s} = 7$ TeV”, *Phys. Rev. D* **85** (2012) [doi:10.1103/physrevd.85.012004](https://doi.org/10.1103/physrevd.85.012004).

[75] CMS Collaboration, “Search for electroweak production of charginos and neutralinos in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2022) 147, [doi:10.1007/JHEP04\(2022\)147](https://doi.org/10.1007/JHEP04(2022)147), [arXiv:2106.14246](https://arxiv.org/abs/2106.14246).

- [76] CMS Collaboration, “CMS SUS Physics Results”. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS>.
- [77] CMS Collaboration, “Search for supersymmetry in proton-proton collisions at 13 TeV in final states with jets and missing transverse momentum”, *JHEP* **10** (2019) 244, [doi:10.1007/JHEP10\(2019\)244](https://doi.org/10.1007/JHEP10(2019)244), [arXiv:1908.04722](https://arxiv.org/abs/1908.04722).
- [78] CMS Collaboration, “Search for supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV with 137 fb^{-1} in final states with a single lepton using the sum of masses of large-radius jets”, *Phys. Rev. D* **101** (2020), no. 5, 052010, [doi:10.1103/PhysRevD.101.052010](https://doi.org/10.1103/PhysRevD.101.052010), [arXiv:1911.07558](https://arxiv.org/abs/1911.07558).
- [79] CMS Collaboration, “Search for direct top squark pair production in events with one lepton, jets, and missing transverse momentum at 13 TeV with the CMS experiment”, *JHEP* **05** (2020) 032, [doi:10.1007/JHEP05\(2020\)032](https://doi.org/10.1007/JHEP05(2020)032), [arXiv:1912.08887](https://arxiv.org/abs/1912.08887).
- [80] CMS Collaboration, “Search for physics beyond the standard model in events with jets and two same-sign or at least three charged leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **80** (2020), no. 8, 752, [doi:10.1140/epjc/s10052-020-8168-3](https://doi.org/10.1140/epjc/s10052-020-8168-3), [arXiv:2001.10086](https://arxiv.org/abs/2001.10086).
- [81] CMS Collaboration, “Search for supersymmetry in proton-proton collisions at $\sqrt{s} = 13$ TeV in events with high-momentum Z bosons and missing transverse momentum”, *JHEP* **09** (2020) 149, [doi:10.1007/JHEP09\(2020\)149](https://doi.org/10.1007/JHEP09(2020)149), [arXiv:2008.04422](https://arxiv.org/abs/2008.04422).
- [82] CMS Collaboration, “Search for top squark pair production using dilepton final states in pp collision data collected at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **81** (2021), no. 1, 3, [doi:10.1140/epjc/s10052-020-08701-5](https://doi.org/10.1140/epjc/s10052-020-08701-5), [arXiv:2008.05936](https://arxiv.org/abs/2008.05936).
- [83] CMS Collaboration, “Search for supersymmetry in final states with two oppositely charged same-flavor leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2021) 123, [doi:10.1007/JHEP04\(2021\)123](https://doi.org/10.1007/JHEP04(2021)123), [arXiv:2012.08600](https://arxiv.org/abs/2012.08600).
- [84] CMS Collaboration, “Search for top squarks in final states with two top quarks and several light-flavor jets in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **104** (2021), no. 3, 032006, [doi:10.1103/PhysRevD.104.032006](https://doi.org/10.1103/PhysRevD.104.032006), [arXiv:2102.06976](https://arxiv.org/abs/2102.06976).
- [85] CMS Collaboration, “Search for top squark production in fully-hadronic final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D*

Bibliography

104 (2021), no. 5, 052001, [doi:10.1103/PhysRevD.104.052001](https://doi.org/10.1103/PhysRevD.104.052001),
[arXiv:2103.01290](https://arxiv.org/abs/2103.01290).

[86] CMS Collaboration, “Search for electroweak production of charginos and neutralinos in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2022) 147, [doi:10.1007/JHEP04\(2022\)147](https://doi.org/10.1007/JHEP04(2022)147), [arXiv:2106.14246](https://arxiv.org/abs/2106.14246).

[87] CMS Collaboration, “Combined searches for the production of supersymmetric top quark partners in proton–proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **81** (2021), no. 11, 970, [doi:10.1140/epjc/s10052-021-09721-5](https://doi.org/10.1140/epjc/s10052-021-09721-5), [arXiv:2107.10892](https://arxiv.org/abs/2107.10892).

[88] CMS Collaboration, “Search for chargino-neutralino production in events with Higgs and W bosons using 137 fb^{-1} of proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **10** (2021) 045, [doi:10.1007/JHEP10\(2021\)045](https://doi.org/10.1007/JHEP10(2021)045), [arXiv:2107.12553](https://arxiv.org/abs/2107.12553).

[89] CMS Collaboration, “Search for supersymmetry in final states with two or three soft leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2022) 091, [doi:10.1007/JHEP04\(2022\)091](https://doi.org/10.1007/JHEP04(2022)091), [arXiv:2111.06296](https://arxiv.org/abs/2111.06296).

[90] CMS Collaboration, “Search for higgsinos decaying to two Higgs bosons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **05** (2022) 014, [doi:10.1007/JHEP05\(2022\)014](https://doi.org/10.1007/JHEP05(2022)014), [arXiv:2201.04206](https://arxiv.org/abs/2201.04206).

[91] CMS Collaboration, “Search for electroweak production of charginos and neutralinos at $\sqrt{s}=13$ TeV in final states containing hadronic decays of WW, WZ, or WH and missing transverse momentum”, *Phys. Lett. B* **842** (2023) 137460, [doi:10.1016/j.physletb.2022.137460](https://doi.org/10.1016/j.physletb.2022.137460), [arXiv:2205.09597](https://arxiv.org/abs/2205.09597).

[92] CMS Collaboration, “Search for direct pair production of supersymmetric partners of τ leptons in the final state with two hadronically decaying τ leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **108** (2023), no. 1, 012011, [doi:10.1103/PhysRevD.108.012011](https://doi.org/10.1103/PhysRevD.108.012011), [arXiv:2207.02254](https://arxiv.org/abs/2207.02254).

[93] CMS Collaboration, “Search for supersymmetry in final states with a single electron or muon using angular correlations and heavy-object identification in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **09** (2023) 149, [doi:10.1007/JHEP09\(2023\)149](https://doi.org/10.1007/JHEP09(2023)149), [arXiv:2211.08476](https://arxiv.org/abs/2211.08476).

- [94] CMS Collaboration, “Search for top squarks in the four-body decay mode with single lepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **06** (2023) 060, [doi:10.1007/JHEP06\(2023\)060](https://doi.org/10.1007/JHEP06(2023)060), [arXiv:2301.08096](https://arxiv.org/abs/2301.08096).
- [95] CMS Collaboration, “Search for top squark pair production in a final state with at least one hadronically decaying tau lepton in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **07** (2023) 110, [doi:10.1007/JHEP07\(2023\)110](https://doi.org/10.1007/JHEP07(2023)110), [arXiv:2304.07174](https://arxiv.org/abs/2304.07174).
- [96] CMS Collaboration, “Search for new physics in multijet events with at least one photon and large missing transverse momentum in proton-proton collisions at 13 TeV”, *JHEP* **10** (2023) 046, [doi:10.1007/JHEP10\(2023\)046](https://doi.org/10.1007/JHEP10(2023)046), [arXiv:2307.16216](https://arxiv.org/abs/2307.16216).
- [97] CMS Collaboration, “Search for supersymmetry in final states with disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV”, [arXiv:2309.16823](https://arxiv.org/abs/2309.16823).
- [98] CMS Collaboration, “Search for stealth supersymmetry in final states with two photons, jets, and low missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, [arXiv:2310.03154](https://arxiv.org/abs/2310.03154).
- [99] CMS Collaboration, “Search for supersymmetry in final states with two oppositely charged same-flavor leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **2021** (2021) [doi:10.1007/jhep04\(2021\)123](https://doi.org/10.1007/jhep04(2021)123).
- [100] H. Baer, V. Barger, and P. Huang, “Hidden SUSY at the LHC: the light higgsino-world scenario and the role of a lepton collider”, *JHEP* **2011** (2011) [doi:10.1007/jhep11\(2011\)031](https://doi.org/10.1007/jhep11(2011)031).
- [101] Z. Han, G. D. Kribs, A. Martin, and A. Menon, “Hunting quasidegenerate higgsinos”, *Phys. Rev. D* **89** (2014) [doi:10.1103/physrevd.89.075007](https://doi.org/10.1103/physrevd.89.075007).
- [102] K. Griest and D. Seckel, “Three exceptions in the calculation of relic abundances”, *Phys. Rev. D* **43** (1991) 3191, [doi:10.1103/PhysRevD.43.3191](https://doi.org/10.1103/PhysRevD.43.3191).
- [103] J. Edsjö and P. Gondolo, “Neutralino relic density including coannihilations”, *Phys. Rev. D* **56** (1997) 1879, [doi:10.1103/physrevd.56.1879](https://doi.org/10.1103/physrevd.56.1879).
- [104] P. Jackson, C. Rogan, and M. Santoni, “Sparticles in motion: Analyzing compressed SUSY scenarios with a new method of event reconstruction”, *Phys. Rev. D* **95** (2017) [doi:10.1103/physrevd.95.035031](https://doi.org/10.1103/physrevd.95.035031).

Bibliography

- [105] L. Lee, C. Ohm, A. Soffer, and T.-T. Yu, “Collider searches for long-lived particles beyond the standard model”, *Progress in Particle and Nuclear Physics* **106** (2019) 210, [doi:10.1016/j.ppnp.2019.02.006](https://doi.org/10.1016/j.ppnp.2019.02.006).
- [106] G. F. Giudice, R. Rattazzi, M. A. Luty, and H. Murayama, “Gaugino mass without singlets”, *JHEP* **1998** (1998) 027, [doi:10.1088/1126-6708/1998/12/027](https://doi.org/10.1088/1126-6708/1998/12/027).
- [107] J. L. Hewett, B. Lillie, M. Masip, and T. G. Rizzo, “Signatures of long-lived gluinos in split supersymmetry”, *JHEP* **2004** (2004) 070, [doi:10.1088/1126-6708/2004/09/070](https://doi.org/10.1088/1126-6708/2004/09/070).
- [108] N. Arkani-Hamed, A. Delgado, and G. Giudice, “The well-tempered neutralino”, *Nucl. Phys. B* **741** (2006) 108, [doi:10.1016/j.nuclphysb.2006.02.010](https://doi.org/10.1016/j.nuclphysb.2006.02.010).
- [109] K. Rolbiecki and K. Sakurai, “Long-lived bino and wino in supersymmetry with heavy scalars and higgsinos”, *JHEP* **2015** (2015) [doi:10.1007/jhep11\(2015\)091](https://doi.org/10.1007/jhep11(2015)091).
- [110] N. Nagata, H. Otono, and S. Shirai, “Probing Bino-Wino Coannihilation at the LHC”, *JHEP* **10** (2015) 086, [doi:10.1007/JHEP10\(2015\)086](https://doi.org/10.1007/JHEP10(2015)086), [arXiv:1506.08206](https://arxiv.org/abs/1506.08206).
- [111] O. S. Brüning et al., “LHC Design Report”. CERN Yellow Reports: Monographs. CERN, Geneva, 2004. [doi:10.5170/CERN-2004-003-V-1](https://doi.org/10.5170/CERN-2004-003-V-1).
- [112] CERN, “LEP design report”, *Report CERN-LEP-84-01*, 1984.
- [113] E. Lopinska, “The CERN accelerator complex, layout in 2022”.
<https://cds.cern.ch/record/2800984>.
- [114] CMS Collaboration, “Public CMS Luminosity Information”. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>.
- [115] T. Sakuma and T. McCauley, “Detector and Event Visualization with SketchUp at the CMS Experiment”, *J. Phys. Conf. Ser.* **513** (2014) 022032, [doi:10.1088/1742-6596/513/2/022032](https://doi.org/10.1088/1742-6596/513/2/022032), [arXiv:1311.4942](https://arxiv.org/abs/1311.4942).
- [116] CMS Collaboration, “The CMS Experiment at the CERN LHC”, *JINST* **3** (2008) S08004, [doi:10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004).
- [117] University of Zurich, CMS Wiki Pages, “3D axis with spherical coordinates & CMS coordinate system”.
<https://wiki.physik.uzh.ch/cms/latex:tikz>.
- [118] CMS Collaboration, “The CMS magnet project: Technical Design Report”, *CMS Technical Design Report CERN-LHCC-97-010, CMS-TDR-1*, 1997.

- [119] CMS Collaboration, “Precise mapping of the magnetic field in the cms barrel yoke using cosmic rays”, *JINST* **5** (2010) T03021, [doi:10.1088/1748-0221/5/03/T03021](https://doi.org/10.1088/1748-0221/5/03/T03021).
- [120] CMS Collaboration, “The CMS tracker system project: Technical Design Report”, [CMS Technical Design Report CERN-LHCC-98-006, CMS-TDR-5](#), 1997.
- [121] CMS Collaboration, “The Phase-2 Upgrade of the CMS Tracker”, [CMS Technical Design Report CERN-LHCC-2017-009, CMS-TDR-014](#), 2017.
- [122] CMS Collaboration, “CMS Technical Design Report for the Pixel Detector Upgrade”, [CMS Technical Design Report CERN-LHCC-2012-016, CMS-TDR-11](#), 2012.
- [123] CMS Collaboration, “The CMS Phase-1 Pixel Detector Upgrade”, *JINST* **16** (2021), no. 02, P02027, [doi:10.1088/1748-0221/16/02/P02027](https://doi.org/10.1088/1748-0221/16/02/P02027), [arXiv:2012.14304](https://arxiv.org/abs/2012.14304).
- [124] CMS Collaboration, “The CMS electromagnetic calorimeter project: Technical Design Report”, [CMS Technical Design Report CERN-LHCC-97-033, CMS-TDR-4](#), 1997.
- [125] A. Benaglia, “The CMS ECAL performance with examples”, *JINST* **9** (2014) C02008, [doi:10.1088/1748-0221/9/02/C02008](https://doi.org/10.1088/1748-0221/9/02/C02008).
- [126] CMS Collaboration, “ECAL 2016 refined calibration and Run 2 summary plots”, [CMS Detector Performance Note CMS-DP-2020-021](#), 2020.
- [127] CMS Collaboration, “The CMS hadron calorimeter project: Technical Design Report”, [CMS Technical Design Report CERN-LHCC-97-031, CMS-TDR-2](#), 1997.
- [128] CMS Collaboration, “Calibration of the CMS hadron calorimeters using proton-proton collision data at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020) P05002, [doi:10.1088/1748-0221/15/05/p05002](https://doi.org/10.1088/1748-0221/15/05/p05002).
- [129] CMS Collaboration, “The CMS muon project: Technical Design Report”, [CMS Technical Design Report CERN-LHCC-97-032, CMS-TDR-3](#), 1997.
- [130] 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](https://doi.org/10.1088/1748-0221/13/06/P06015), [arXiv:1804.04528](https://arxiv.org/abs/1804.04528).
- [131] A. Colaleo, A. Safonov, A. Sharma, and M. Tytgat, “CMS Technical Design Report for the Muon Endcap GEM Upgrade”, [CMS Technical](#)

Bibliography

Design Report CERN-LHCC-2015-012, CMS-TDR-013, 2015.

[132] CMS Collaboration, “CMS TriDAS project: Technical Design Report, Volume 1: The Trigger Systems”, [CMS Technical Design Report CERN-LHCC-2000-038, CMS-TDR-6-1](#), 2000.

[133] CMS Collaboration, “CMS Technical Design Report for the Level-1 Trigger Upgrade”, [CMS Technical Design Report CERN-LHCC-2013-011, CMS-TDR-12](#), 2013.

[134] W. Adam et al., “The CMS high level trigger”, *Eur. Phys. J. C* **46** (2006) 605–667, [doi:10.1140/epjc/s2006-02495-8](https://doi.org/10.1140/epjc/s2006-02495-8).

[135] CMS Collaboration, “CMS: The computing project. Technical design report”, [CMS Technical Design Report CERN-LHCC-2005-023, CMS-TDR-7](#), 6, 2005.

[136] CMS Collaboration, “Data Parking and Data Scouting at the CMS Experiment”, [CMS Detector Performance Note CMS-DP-2012-022](#), 2012.

[137] CMS Collaboration, “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020), no. 10, P10017, [doi:10.1088/1748-0221/15/10/P10017](https://doi.org/10.1088/1748-0221/15/10/P10017), [arXiv:2006.10165](https://arxiv.org/abs/2006.10165).

[138] CMS Collaboration, “The CMS calorimeter trigger for LHC Run II”, *J. Phys.: Conf. Ser.* **928** (2017), no. 1, 012004, [doi:10.1088/1742-6596/928/1/012004](https://doi.org/10.1088/1742-6596/928/1/012004).

[139] CMS Collaboration, “Run 2 Upgrades to the CMS Level-1 Calorimeter Trigger”, *JINST* **11** (2016), no. 01, C01051, [doi:10.1088/1748-0221/11/01/C01051](https://doi.org/10.1088/1748-0221/11/01/C01051), [arXiv:1511.05855](https://arxiv.org/abs/1511.05855).

[140] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_t jet clustering algorithm”, *JHEP* **2008** (2008) 063, [doi:10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063).

[141] CMS Collaboration, “Performance of the CMS muon trigger system in proton-proton collisions at $\sqrt{s} = 13$ ”, *JINST* **16** (2021) P07001, [doi:10.1088/1748-0221/16/07/p07001](https://doi.org/10.1088/1748-0221/16/07/p07001).

[142] I. Papavergou, “The CMS Level-1 muon triggers for the LHC Run II”, [doi:10.22323/1.321.0070](https://doi.org/10.22323/1.321.0070).

[143] O. Aberle et al., “High-Luminosity Large Hadron Collider (HL-LHC): Technical design report”, [Technical Design Report CERN-2020-010](#), 2020.

- [144] D. Contardo et al., “Technical Proposal for the Phase-II Upgrade of the CMS Detector”, [doi:10.17181/CERN.VU8I.D59J](https://doi.org/10.17181/CERN.VU8I.D59J).
- [145] CMS Collaboration, “The Phase-2 Upgrade of the CMS Endcap Calorimeter”, [CMS Technical Design Report CERN-LHCC-2017-023, CMS-TDR-019](https://cds.cern.ch/record/2234413), 2017.
- [146] CMS Collaboration, “The Phase-2 Upgrade of the CMS Barrel Calorimeters”, [CMS Technical Design Report CERN-LHCC-2017-011, CMS-TDR-015](https://cds.cern.ch/record/2234414), 2017.
- [147] CMS Collaboration, “The Phase-2 Upgrade of the CMS Muon Detectors”, [CMS Technical Design Report CERN-LHCC-2017-012, CMS-TDR-016](https://cds.cern.ch/record/2234415), 2017.
- [148] CMS Collaboration, “The Phase-2 Upgrade of the CMS Level-1 Trigger”, [CMS Technical Design Report CERN-LHCC-2020-004, CMS-TDR-021](https://cds.cern.ch/record/2234416), 2020.
- [149] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017) P10003, [doi:10.1088/1748-0221/12/10/P10003](https://doi.org/10.1088/1748-0221/12/10/P10003), [arXiv:1706.04965](https://arxiv.org/abs/1706.04965).
- [150] CMS Collaboration, “Description and performance of track and primary-vertex reconstruction with the CMS tracker”, *JINST* **9** (2014) P10009, [doi:10.1088/1748-0221/9/10/p10009](https://doi.org/10.1088/1748-0221/9/10/p10009).
- [151] R. Frühwirth, “Application of kalman filtering to track and vertex fitting”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **262** (1987) 444, [doi:10.1016/0168-9002\(87\)90887-4](https://doi.org/10.1016/0168-9002(87)90887-4).
- [152] P. Billoir, “Progressive track recognition with a kalman-like fitting procedure”, *Computer Physics Communications* **57** (1989) 390, [doi:10.1016/0010-4655\(89\)90249-X](https://doi.org/10.1016/0010-4655(89)90249-X).
- [153] P. Billoir and S. Qian, “Simultaneous pattern recognition and track fitting by the kalman filtering method”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **294** (1990) 219, [doi:10.1016/0168-9002\(90\)91835-Y](https://doi.org/10.1016/0168-9002(90)91835-Y).
- [154] S. Cucciarelli, M. Konecki, D. Kotlinski, and T. Todorov, “Track reconstruction, primary vertex finding and seed generation with the Pixel Detector”, [CMS Note CMS-NOTE-2006-026](https://cds.cern.ch/record/2006-026), 2006.

Bibliography

- [155] C. Amsler et al., “Review of particle physics”, *Phys. Lett. B* **667** (2008) 1, [doi:10.1016/j.physletb.2008.07.018](https://doi.org/10.1016/j.physletb.2008.07.018).
- [156] H. Bethe and W. Heitler, “On the Stopping of fast particles and on the creation of positive electrons”, *Proc. Roy. Soc. Lond. A* **146** (1934) 83–112, [doi:10.1098/rspa.1934.0140](https://doi.org/10.1098/rspa.1934.0140).
- [157] R. Fröhwirth, “A gaussian-mixture approximation of the Bethe–Heitler model of electron energy loss by bremsstrahlung”, *Computer Physics Communications* **154** (2003) 131, [doi:10.1016/S0010-4655\(03\)00292-3](https://doi.org/10.1016/S0010-4655(03)00292-3).
- [158] CMS Collaboration, “2017 tracking performance plots”, [CMS Detector Performance Note CMS-DP-2017-015](#), 2017.
- [159] CMS Collaboration, “Muon Reconstruction and Identification Improvements for Run-2 and First Results with 2015 Run Data”, [CMS Detector Performance Note CMS-DP-2015-015](#), 2015.
- [160] F. Pantaleo, “New Track Seeding Techniques for the CMS Experiment”, [Thesis CMS-TS-2017-028](#), 2017.
- [161] K. Rose, “Deterministic annealing for clustering, compression, classification, regression, and related optimization problems”, *Proceedings of the IEEE* **86** (1998), no. 11, 2210–2239, [doi:10.1109/5.726788](https://doi.org/10.1109/5.726788).
- [162] R. Fröhwirth, W. Waltenberger, and P. Vanlaer, “Adaptive Vertex Fitting”, [CMS Note CMS-NOTE-2007-008](#), 2007.
- [163] M. Cacciari, G. P. Salam, and G. Soyez, “FastJet user manual”, *The European Physical Journal C* **72** (2012) [doi:10.1140/epjc/s10052-012-1896-2](https://doi.org/10.1140/epjc/s10052-012-1896-2).
- [164] CMS Collaboration, “Pileup mitigation at CMS in 13 TeV data”, *JINST* **15** (2020), no. 09, P09018, [doi:10.1088/1748-0221/15/09/P09018](https://doi.org/10.1088/1748-0221/15/09/P09018), [arXiv:2003.00503](https://arxiv.org/abs/2003.00503).
- [165] CMS Collaboration, “Studies of Tracker Material”, [CMS Physics Analysis Summary CMS-PAS-TRK-10-003](#), 2010.
- [166] CMS Collaboration, “CMS Tracking Performance Results from Early LHC Operation”, *Eur. Phys. J. C* **70** (2010) 1165, [doi:10.1140/epjc/s10052-010-1491-3](https://doi.org/10.1140/epjc/s10052-010-1491-3), [arXiv:1007.1988](https://arxiv.org/abs/1007.1988).
- [167] CMS Collaboration, “Performance of photon reconstruction and identification with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *JINST* **10** (2015) P08010, [doi:10.1088/1748-0221/10/08/p08010](https://doi.org/10.1088/1748-0221/10/08/p08010).

- [168] CMS Collaboration, “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV”, *JINST* **13** (2018) P05011, [doi:10.1088/1748-0221/13/05/P05011](https://doi.org/10.1088/1748-0221/13/05/P05011), [arXiv:1712.07158](https://arxiv.org/abs/1712.07158).
- [169] CMS Collaboration, “CMS Physics: Technical Design Report Volume 1: Detector Performance and Software”, [Technical Report CERN-LHCC-2006-001, CMS-TDR-8-1](https://cds.cern.ch/record/1013001), 2006.
- [170] CMS Collaboration, “Performance of CMS muon reconstruction in cosmic-ray events”, *JINST* **5** (2010) T03022, [doi:10.1088/1748-0221/5/03/t03022](https://doi.org/10.1088/1748-0221/5/03/t03022).
- [171] CMS Collaboration, “Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV”, *JINST* **7** (2012) P10002, [doi:10.1088/1748-0221/7/10/P10002](https://doi.org/10.1088/1748-0221/7/10/P10002).
- [172] CMS Collaboration, “Particle-flow commissioning with muons and electrons from J/Psi and W events at 7 TeV”, [CMS Physics Analysis Summary CMS-PAS-PFT-10-003](https://cds.cern.ch/record/1013003), 2010.
- [173] CMS Collaboration, “Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *JINST* **10** (2015) P06005, [doi:10.1088/1748-0221/10/06/p06005](https://doi.org/10.1088/1748-0221/10/06/p06005).
- [174] CMS Collaboration, “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”, *JINST* **16** (2021) P05014, [doi:10.1088/1748-0221/16/05/p05014](https://doi.org/10.1088/1748-0221/16/05/p05014).
- [175] CMS Collaboration, “Recording and reconstructing 10 billion unbiased b hadron decays in CMS”, [CMS Detector Performance Note CMS-DP-2019-043](https://cds.cern.ch/record/2604303), 2019.
- [176] CMS Collaboration, “Recording and reconstructing 10 billion unbiased b hadron decays in CMS”, [doi:10.1051/epjconf/202024501025](https://doi.org/10.1051/epjconf/202024501025).
- [177] 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** (2005) N9, [doi:10.1088/0954-3899/31/9/n01](https://doi.org/10.1088/0954-3899/31/9/n01).
- [178] A. Hoecker et al., “TMVA - toolkit for multivariate data analysis”, [arXiv:physics/0703039](https://arxiv.org/abs/physics/0703039).
- [179] CracowAFPGroup Collaboration, “Study of Diffractive Bremsstrahlung at 13 TeV LHC”, *Acta Phys. Pol. B* **47** (2016) 1639–1644, [doi:10.5506/APhysPolB.47.1639](https://doi.org/10.5506/APhysPolB.47.1639).

Bibliography

- [180] N. Marinelli, “Track finding and identification of converted photons”, [CMS Note CMS-NOTE-2006-005](#), 2006.
- [181] CMS Collaboration, “Photon reconstruction and identification at $\sqrt{s} = 7$ TeV”, [CMS Physics Analysis Summary CMS-PAS-EGM-10-005](#), 2010.
- [182] LHCb Collaboration, “Search for lepton-universality violation in $B^+ \rightarrow K^+ \ell^+ \ell^-$ decays”, *Phys. Rev. Lett.* **122** (2019) 191801, [doi:10.1103/PhysRevLett.122.191801](https://doi.org/10.1103/PhysRevLett.122.191801).
- [183] LHCb Collaboration, “Test of lepton universality with $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays”, *JHEP* **08** (2017) 055, [doi:10.1007/JHEP08\(2017\)055](https://doi.org/10.1007/JHEP08(2017)055), [arXiv:1705.05802](https://arxiv.org/abs/1705.05802).
- [184] Belle Collaboration, “Test of Lepton-Flavor Universality in $B \rightarrow K^* \ell^+ \ell^-$ Decays at Belle”, *Phys. Rev. Lett.* **126** (2021) 161801, [doi:10.1103/PhysRevLett.126.161801](https://doi.org/10.1103/PhysRevLett.126.161801), [arXiv:1904.02440](https://arxiv.org/abs/1904.02440).
- [185] BABAR Collaboration, “Evidence for an excess of $\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau$ decays”, *Phys. Rev. Lett.* **109** (2012) 101802, [doi:10.1103/PhysRevLett.109.101802](https://doi.org/10.1103/PhysRevLett.109.101802).
- [186] Belle Collaboration, “Measurement of the branching ratio of $\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau$ relative to $\bar{B} \rightarrow D^{(*)} \ell^- \bar{\nu}_\ell$ decays with hadronic tagging at belle”, *Phys. Rev. D* **92** (2015) 072014, [doi:10.1103/PhysRevD.92.072014](https://doi.org/10.1103/PhysRevD.92.072014).
- [187] LHCb Collaboration, “Test of lepton flavor universality by the measurement of the $B^0 \rightarrow D^{*-} \tau^+ \nu_\tau$ branching fraction using three-prong τ decays”, *Phys. Rev. D* **97** (2018) 072013, [doi:10.1103/PhysRevD.97.072013](https://doi.org/10.1103/PhysRevD.97.072013).
- [188] LHCb Collaboration, “Measurement of CP -Averaged Observables in the $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ Decay”, *Phys. Rev. Lett.* **125** (2020) 011802, [doi:10.1103/PhysRevLett.125.011802](https://doi.org/10.1103/PhysRevLett.125.011802), [arXiv:2003.04831](https://arxiv.org/abs/2003.04831).
- [189] Belle Collaboration, “Angular analysis of $B^0 \rightarrow K^*(892)^0 \ell^+ \ell^-$ ”, [Belle Conference Paper BELLE-CONF-1603](#), 2016.
- [190] BaBar Collaboration, “Measurement of angular asymmetries in the decays $b \rightarrow K^* \ell^+ \ell^-$ ”, *Phys. Rev. D* **93** (2016) 052015, [doi:10.1103/PhysRevD.93.052015](https://doi.org/10.1103/PhysRevD.93.052015).
- [191] V. Khachatryan et al., “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”, *JINST* **12** (2017) P02014, [doi:10.1088/1748-0221/12/02/p02014](https://doi.org/10.1088/1748-0221/12/02/p02014).

- [192] CMS Collaboration, “Jet energy scale and resolution performance with 13 TeV data collected by CMS in 2016-2018”, [CMS Detector Performance Note CMS-DP-2020-019](#), 2020.
- [193] CMS Collaboration, “Track impact parameter resolution for the full pseudo rapidity coverage in the 2017 dataset with the CMS Phase-1 Pixel detector”, [CMS Detector Performance Note CMS-DP-2020-049](#), 2020.
- [194] CMS Collaboration, “Identification of b-quark jets with the CMS experiment”, *JINST* **8** (2013) P04013, [doi:10.1088/1748-0221/8/04/p04013](https://doi.org/10.1088/1748-0221/8/04/p04013).
- [195] D. Guest et al., “Jet Flavor Classification in High-Energy Physics with Deep Neural Networks”, *Phys. Rev. D* **94** (2016) 112002, [doi:10.1103/PhysRevD.94.112002](https://doi.org/10.1103/PhysRevD.94.112002), [arXiv:1607.08633](https://arxiv.org/abs/1607.08633).
- [196] E. Bols et al., “Jet flavour classification using DeepJet”, *JINST* **15** (2020) P12012, [doi:10.1088/1748-0221/15/12/p12012](https://doi.org/10.1088/1748-0221/15/12/p12012).
- [197] CMS Collaboration, “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018) P10005, [doi:10.1088/1748-0221/13/10/p10005](https://doi.org/10.1088/1748-0221/13/10/p10005).
- [198] CMS Collaboration, “Performance of τ -lepton reconstruction and identification in CMS”, *JINST* **7** (2012) P01001, [doi:10.1088/1748-0221/7/01/p01001](https://doi.org/10.1088/1748-0221/7/01/p01001).
- [199] CMS Collaboration, “Reconstruction and identification of τ lepton decays to hadrons and ν_τ at CMS”, *JINST* **11** (2016) P01019, [doi:10.1088/1748-0221/11/01/p01019](https://doi.org/10.1088/1748-0221/11/01/p01019).
- [200] CMS Collaboration, “Performance of the DeepTau algorithm for the discrimination of taus against jets, electron, and muons”, [CMS Detector Performance Note CMS-DP-2019-033](#), 2019.
- [201] CMS Collaboration, “Performance of the CMS missing transverse momentum reconstruction in pp data at $\sqrt{s} = 8$ TeV”, *JINST* **10** (2015) P02006, [doi:10.1088/1748-0221/10/02/P02006](https://doi.org/10.1088/1748-0221/10/02/P02006), [arXiv:1411.0511](https://arxiv.org/abs/1411.0511).
- [202] CMS Collaboration, “Performance of missing transverse momentum reconstruction in events containing a photon and jets collected by CMS during proton-proton collisions at $\sqrt{s} = 13$ TeV in 2018”, [CMS Detector Performance Note CMS-DP-2020-031](#), 2020.
- [203] D. Bertolini et al., “Pileup per particle identification”, *JHEP* **10** (2014) 059, [doi:10.1007/jhep10\(2014\)059](https://doi.org/10.1007/jhep10(2014)059), [arXiv:1407.6013](https://arxiv.org/abs/1407.6013).

Bibliography

- [204] S. Jung, “Resolving the existence of higgsinos in the LHC inverse problem”, *JHEP* **2014** (2014), no. 6, [doi:10.1007/jhep06\(2014\)111](https://doi.org/10.1007/jhep06(2014)111).
- [205] CMS Collaboration, “Search for new physics in events with two soft oppositely charged leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **782** (2018) 440–467, [doi:10.1016/j.physletb.2018.05.062](https://doi.org/10.1016/j.physletb.2018.05.062), [arXiv:1801.01846](https://arxiv.org/abs/1801.01846).
- [206] CMS Collaboration, “Search for supersymmetry in events with soft leptons, low jet multiplicity, and missing transverse energy in proton–proton collisions at $\sqrt{s}=8$ TeV”, *Phys. Lett. B* **759** (2016) 9–35, [doi:10.1016/j.physletb.2016.05.033](https://doi.org/10.1016/j.physletb.2016.05.033), [arXiv:1512.08002](https://arxiv.org/abs/1512.08002).
- [207] ATLAS Collaboration, “Search for direct production of charginos and neutralinos in events with three leptons and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector”, *JHEP* **04** (2014) 169, [doi:10.1007/JHEP04\(2014\)169](https://doi.org/10.1007/JHEP04(2014)169), [arXiv:1402.7029](https://arxiv.org/abs/1402.7029).
- [208] ATLAS Collaboration, “Search for electroweak production of supersymmetric states in scenarios with compressed mass spectra at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Phys. Rev. D* **97** (2018), no. 5, 052010, [doi:10.1103/PhysRevD.97.052010](https://doi.org/10.1103/PhysRevD.97.052010), [arXiv:1712.08119](https://arxiv.org/abs/1712.08119).
- [209] ATLAS Collaboration, “Searches for electroweak production of supersymmetric particles with compressed mass spectra in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”, *Phys. Rev. D* **101** (2020), no. 5, 052005, [doi:10.1103/PhysRevD.101.052005](https://doi.org/10.1103/PhysRevD.101.052005), [arXiv:1911.12606](https://arxiv.org/abs/1911.12606).
- [210] ATLAS Collaboration, “Search for chargino–neutralino pair production in final states with three leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”, *Eur. Phys. J. C* **81** (2021), no. 12, 1118, [doi:10.1140/epjc/s10052-021-09749-7](https://doi.org/10.1140/epjc/s10052-021-09749-7), [arXiv:2106.01676](https://arxiv.org/abs/2106.01676).
- [211] U. De Sanctis, T. Lari, S. Montesano, and C. Troncon, “Perspectives for the detection and measurement of Supersymmetry in the focus point region of mSUGRA models with the ATLAS detector at LHC”, *Eur. Phys. J. C* **52** (2007) 743–758, [doi:10.1140/epjc/s10052-007-0415-3](https://doi.org/10.1140/epjc/s10052-007-0415-3), [arXiv:0704.2515](https://arxiv.org/abs/0704.2515).
- [212] T. Sjöstrand et al., “An Introduction to PYTHIA 8.2”, *Comput. Phys. Commun.* **191** (2015) 159, [doi:10.1016/j.cpc.2015.01.024](https://doi.org/10.1016/j.cpc.2015.01.024), [arXiv:1410.3012](https://arxiv.org/abs/1410.3012).

[213] A. Djouadi, M. M. Muhlleitner, and M. Spira, “Decays of supersymmetric particles: the program susy-hit (suspect-sdecay-hdecay-interface)”, 2006.

[214] CMS Collaboration, “Combined search for electroweak production of winos, binos, higgsinos, and sleptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **109** (2024) 112001, [doi:10.1103/PhysRevD.109.112001](https://doi.org/10.1103/PhysRevD.109.112001), [arXiv:2402.01888](https://arxiv.org/abs/2402.01888).

[215] B. Allanach, “Softsusy: A program for calculating supersymmetric spectra”, *Computer Physics Communications* **143** (2002) 305–331, [doi:10.1016/s0010-4655\(01\)00460-x](https://doi.org/10.1016/s0010-4655(01)00460-x).

[216] J. Alwall et al., “MadGraph5: going beyond”, *JHEP* **06** (2011) 128, [doi:10.1007/JHEP06\(2011\)128](https://doi.org/10.1007/JHEP06(2011)128), [arXiv:1106.0522](https://arxiv.org/abs/1106.0522).

[217] 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](https://doi.org/10.1007/JHEP07(2014)079), [arXiv:1405.0301](https://arxiv.org/abs/1405.0301).

[218] 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](https://doi.org/10.1088/1126-6708/2004/11/040), [arXiv:hep-ph/0409146](https://arxiv.org/abs/hep-ph/0409146).

[219] 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](https://doi.org/10.1088/1126-6708/2007/11/070), [arXiv:0709.2092](https://arxiv.org/abs/0709.2092).

[220] 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](https://doi.org/10.1007/JHEP06(2010)043), [arXiv:1002.2581](https://arxiv.org/abs/1002.2581).

[221] S. Alioli, P. Nason, C. Oleari, and E. Re, “NLO single-top production matched with shower in POWHEG: s - and t -channel contributions”, *JHEP* **09** (2009) 111, [doi:10.1088/1126-6708/2009/09/111](https://doi.org/10.1088/1126-6708/2009/09/111), [arXiv:0907.4076](https://arxiv.org/abs/0907.4076). [Erratum: *JHEP* **02** (2010) 011].

[222] E. Re, “Single-top Wt -channel production matched with parton showers using the POWHEG method”, *Eur. Phys. J. C* **71** (2011) 1547, [doi:10.1140/epjc/s10052-011-1547-z](https://doi.org/10.1140/epjc/s10052-011-1547-z), [arXiv:1009.2450](https://arxiv.org/abs/1009.2450).

[223] NNPDF Collaboration, “Parton distributions from high-precision collider data”, *Eur. Phys. J. C* **77** (2017) 663, [doi:10.1140/epjc/s10052-017-5199-5](https://doi.org/10.1140/epjc/s10052-017-5199-5).

Bibliography

- [224] 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](https://doi.org/10.1140/epjc/s10052-019-7499-4), [arXiv:1903.12179](https://arxiv.org/abs/1903.12179).
- [225] S. Agostinelli et al., “GEANT4 — a simulation toolkit”, *Nucl. Instr. and Meth. A* **506** (2003) 250, [doi:10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [226] CMS Collaboration, “Search for neutral Higgs bosons decaying to tau pairs in pp collisions at $\sqrt{s}=7$ TeV”, *Physics Letters B* **713** (2012) 68–90, [doi:10.1016/j.physletb.2012.05.028](https://doi.org/10.1016/j.physletb.2012.05.028).
- [227] Z. Han, G. D. Kribs, A. Martin, and A. Menon, “Hunting quasidegenerate higgsinos”, *Physical Review D* **89** (2014), no. 7, [doi:10.1103/physrevd.89.075007](https://doi.org/10.1103/physrevd.89.075007).
- [228] CMS Collaboration, “Performance of Methods for Data-Driven Background Estimation in SUSY Searches”, **CMS Physics Analysis Summary CMS-PAS-SUS-10-001**, 2010.
- [229] T. C. Collaboration, “Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 13$ TeV”, *JHEP* **2018** (2018) 161, [doi:10.1007/JHEP07\(2018\)161](https://doi.org/10.1007/JHEP07(2018)161), [arXiv:1802.02613](https://arxiv.org/abs/1802.02613).
- [230] ALEPH, DELPHI, L3, OPAL Experiments, “Combined LEP Chargino Results, up to 208 for low DM”, **LEPSUSYWG/02-04.1**, 2002.
- [231] ATLAS Collaboration, “Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Eur. Phys. J. C* **82** (2022), no. 7, 606, [doi:10.1140/epjc/s10052-022-10489-5](https://doi.org/10.1140/epjc/s10052-022-10489-5), [arXiv:2201.02472](https://arxiv.org/abs/2201.02472).
- [232] ATLAS Collaboration, “Search for nearly mass-degenerate higgsinos using low-momentum mildly-displaced tracks in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, [arXiv:2401.14046](https://arxiv.org/abs/2401.14046).
- [233] CMS Collaboration, “Search for supersymmetry in final states with two oppositely charged same-flavor leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **04** (2021) 123, [doi:10.1007/JHEP04\(2021\)123](https://doi.org/10.1007/JHEP04(2021)123), [arXiv:2012.08600](https://arxiv.org/abs/2012.08600).
- [234] CMS Collaboration, “Combined search for electroweak production of charginos and neutralinos in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **03** (2018) 160, [doi:10.1007/JHEP03\(2018\)160](https://doi.org/10.1007/JHEP03(2018)160), [arXiv:1801.03957](https://arxiv.org/abs/1801.03957).
- [235] K. Matchev and S. Thomas, “Higgs and Z-boson signatures of supersymmetry”, *Physical Review D* **62** (2000), no. 7,

[doi:10.1103/physrevd.62.077702](https://doi.org/10.1103/physrevd.62.077702).

[236] H. Baer et al., “Natural SUSY with a bino- or wino-like LSP”, *Phys. Rev. D* **91** (2015), no. 7, [doi:10.1103/physrevd.91.075005](https://doi.org/10.1103/physrevd.91.075005).

[237] C. G. Lester and D. J. Summers, “Measuring masses of semiinvisibly decaying particles pair produced at hadron colliders”, *Phys. Lett. B* **463** (1999) 99, [doi:10.1016/S0370-2693\(99\)00945-4](https://doi.org/10.1016/S0370-2693(99)00945-4), [arXiv:hep-ph/9906349](https://arxiv.org/abs/hep-ph/9906349).

[238] A. Barr, C. Lester, and P. Stephens, “m(T2): The truth behind the glamour”, *J. Phys. G* **29** (2003) 2343, [doi:10.1088/0954-3899/29/10/304](https://doi.org/10.1088/0954-3899/29/10/304), [arXiv:hep-ph/0304226](https://arxiv.org/abs/hep-ph/0304226).

[239] CMS Collaboration, “Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques”, *JINST* **15** (2020) P06005, [doi:10.1088/1748-0221/15/06/P06005](https://doi.org/10.1088/1748-0221/15/06/P06005), [arXiv:2004.08262](https://arxiv.org/abs/2004.08262).

[240] ATLAS and CMS Collaborations, LHC Higgs Combination Group, “Procedure for the LHC Higgs boson search combination in Summer 2011”, Technical Report CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-11, 2011.

[241] T. Junk, “Confidence level computation for combining searches with small statistics”, *Nucl. Instrum. Meth. A* **434** (1999) 435, [doi:10.1016/S0168-9002\(99\)00498-2](https://doi.org/10.1016/S0168-9002(99)00498-2), [arXiv:hep-ex/9902006](https://arxiv.org/abs/hep-ex/9902006).

[242] A. L. Read, “Presentation of search results: the CLs technique”, *J. Phys. G* **28** (2002) 2693, [doi:10.1088/0954-3899/28/10/313](https://doi.org/10.1088/0954-3899/28/10/313).

[243] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics”, *Eur. Phys. J. C* **71** (2011) 1554, [doi:10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0), [arXiv:1007.1727](https://arxiv.org/abs/1007.1727).

[244] CMS Collaboration, “Electron Reconstruction and Identification in the CMS Phase-2 Level-1 Trigger”, CMS Detector Performance Note CMS-DP-2023-047, 2023.

[245] J. Duarte et al., “Fast inference of deep neural networks in FPGAs for particle physics”, *JINST* **13** (2018), no. 07, P07027, [doi:10.1088/1748-0221/13/07/p07027](https://doi.org/10.1088/1748-0221/13/07/p07027).

[246] S. Summers et al., “Fast inference of boosted decision trees in FPGAs for particle physics”, *JINST* **15** (2020), no. 05, P05026–P05026, [doi:10.1088/1748-0221/15/05/p05026](https://doi.org/10.1088/1748-0221/15/05/p05026).

[247] E. Bartz et al., “FPGA-based tracking for the CMS Level-1 trigger using the tracklet algorithm”, *JINST* **15** (2020), no. 06, P06024,

Bibliography

[doi:10.1088/1748-0221/15/06/P06024](https://doi.org/10.1088/1748-0221/15/06/P06024), [arXiv:1910.09970](https://arxiv.org/abs/1910.09970).

- [248] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system”, *CoRR* (2016) [arXiv:1603.02754](https://arxiv.org/abs/1603.02754).
- [249] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions”, in *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [250] L. Lista, “Statistical Methods for Data Analysis in Particle Physics”. Springer, 2017. ISBN 978-3-319-20175-7.