

Search for Stealth Supersymmetry in Proton-Proton Collisions with the CMS Experiment at the Large Hadron Collider

by

Tanmay Kamalakar Mudholkar

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

at

Carnegie Mellon University
Department of Physics
Pittsburgh, Pennsylvania

Advised by Professor Manfred Paulini

September 27, 2022



Abstract

This Thesis documents a search for Stealth supersymmetry (SUSY) in final states with two photons, jets, and low missing transverse momentum. Over the data-taking period of the LHC, a progressively larger fraction of the well-motivated parameter-space for typical SUSY models has been ruled out experimentally, but these signal models typically assume a signature with high missing transverse momentum. The study documented here helps close the gap by extending the search to a specific subset of SUSY models called “Stealth SUSY” that do not have this signature, but happen to be nearly as phenomenologically well-motivated as typical minimal supersymmetric Standard Model scenarios. An innovative “hybrid” approach is used to estimate the background, which involves the augmentation of a zeroth-order data-driven background model with first-order corrections obtained from Monte Carlo simulations of the background. The search is based on a 138 fb^{-1} sample of proton-proton collisions collected at a center-of-mass energy of 13 TeV with the CMS detector. The results are interpreted in the context of simplified stealth models with gluino and squark pair production. The observed data are consistent with the Standard Model prediction, and gluino (squark) masses of up to 2.1 (1.85) TeV are excluded.

Acknowledgments

I would like to express my profound gratitude to my advisor, Manfred Paulini, whose constant encouragement and guidance over the years has been critical to the successful conclusion of the search for Stealth SUSY that has become the subject of this Thesis. The object of scientific research is to push the boundaries of the known, and the paths that lead into the unknown are quite unpredictable and nonlinear. This work turned out to be significantly subtler and more demanding than I had anticipated when I started working on it; Manfred's long experience has been invaluable in finding the right balance between chasing new ideas and fine-tuning existing ones. How best to convey a novel technical concept in a document or presentation is an often overlooked skill, and Manfred's many suggestions have made a lasting impression on the way I tackle such communication. I also deeply appreciate all his hard work to ensure that I always had fulfilling projects to work on, and adequate funding to allow me to devote my full attention to my research despite all the delays.

More broadly, this Thesis is only possible due to the combined efforts of the LHC accelerator team and the CMS collaboration. The pursuit of knowledge for its own sake is truly an ideal that transcends cultures and borders, and I am immensely grateful to all the scientists and engineers whose choice to cooperate for the sake of this pursuit, often without the expectation of individual recognition, has led to the present abundance of experimental data that allows us to ask fundamental questions of Nature – and to receive answers to some of them. It has been the honor and privilege of a lifetime to have been given the opportunity to join this pursuit as a member of the CMS collaboration. The ECAL team in particular always made me feel at home at CERN, and I am grateful for their patience in answering all my questions and their willingness to delegate to fresh students some responsibility over the functioning of the detector.

I am grateful for the advice, help, and camaraderie of all members of the growing CMS group at Carnegie Mellon. Michael Andrews has been a constant presence throughout my PhD as a friend and colleague, and has always been a pleasure to work with – in addition to introducing me to the nuts and bolts of DQM and the CMS analysis software, he has always been a good sounding-board for new ideas. John Alison and Patrick Bryant contributed several useful comments and suggestions during our regular group meetings, many of which have influenced this Thesis. Menglei Sun, Alejandro Sanchez, and Abhirami Harilal have all been wonderful colleagues, and their hard work in maintaining and extending DQM continues to provide a valuable service to ECAL.

We all stand on the shoulders of giants. I owe all who have taught me physics an immense debt of gratitude. I have had the opportunity to learn from numerous excellent teachers who have gradually opened my eyes to the beauty and structure of physics, from simple mechanical systems to advanced topics in quantum field theory.

I wish to thank in particular Debashish Chowdhury, Michael Widom, and Pankaj Jain, who guided me in my earliest forays into academic research as an undergraduate student.

I am deeply grateful to all my friends for their support and companionship, and for volunteering as subjects in my culinary trials. Aaditya Ramdas and Devashish Gopalan were wonderful roommates during my first stint in Pittsburgh, and continue to be good friends. Nabarun Dev and Fanbo Meng made my stay at CERN more rewarding with many stimulating conversations. Vikesh Siddhu, Mert Terzi, Leila Wehbe, Yash Puranik and Jessica Cheng have all provided delightful company over many years.

Finally, I would like to express my eternal gratitude to my family. During my crucial formative years, my parents, Kamalakar and Kalpana Mudholkar, ensured that I had the resources to satiate my curiosity, and an environment in which learning about the world was encouraged and rewarded. My sister, Suchita Mudholkar, has always been a constant source of happiness and good advice, and our conversations continue to keep my mind supplied with invaluable reminders of the fact that the world outside these ivory towers is as rich and diverse as the world inside is fulfilling.

Contents

Contents	iii
List of Tables	viii
List of Figures	ix
Organization and Notation	1
1 Theoretical Background	4
1.1 Introduction to the SM	4
1.2 The Mathematical Structure of the SM	6
1.2.1 SM Lagrangian	6
The Lagrangian (before spontaneous symmetry breaking) . . .	6
Gauge sector	6
Lepton sector	8
Quark sector	9
Higgs Sector	10
Yukawa couplings	11
Full SM Lagrangian	11
1.2.2 Spontaneous Symmetry Breaking	12
Consequences of spontaneous symmetry breaking: gauge sector	13
Consequences of spontaneous symmetry breaking: lepton and	
quark sectors	14
1.3 Issues with the SM	15
1.3.1 Hierarchy Problem	15
1.3.2 Running of the Gauge Couplings	17
1.3.3 Neutrino Masses	18
1.3.4 Particle Nature of Dark Matter	18
1.4 Theoretical Underpinnings of SUSY	19
1.4.1 SUSY as an Extension of the SM	19
Chiral Superfields	21
Vector Superfields	22

1.4.2	SUSY Lagrangians	22
	F -term	23
	D -term	23
	SUSY Lagrangian with no gauging	23
	Gauging the SUSY Lagrangian	24
	Lagrangian with full gauge group	27
1.4.3	Mass degeneracy of superpartners in SUSY	27
1.4.4	The Minimal Supersymmetric Standard Model (MSSM)	28
1.4.5	Phenomenological Modifications in the MSSM	30
	Soft SUSY breaking	30
	R -parity	32
1.5	Motivation for SUSY	33
1.5.1	Solution to the Hierarchy Problem	33
1.5.2	Gauge Coupling Unification	34
1.5.3	The WIMP Miracle	34
2	Phenomenology and Stealth SUSY	37
2.1	Current Status of SUSY searches	37
2.2	Stealth SUSY	40
2.2.1	Extensions to the MSSM with low energy imbalance	40
2.2.2	Brief Description of Stealth SUSY	41
2.2.3	Motivation for Stealth SUSY	42
2.2.4	Previous Searches for Stealth SUSY	44
3	Detector Description	46
3.1	The LHC	46
3.1.1	Overview of the accelerator complex	46
3.1.2	Accelerator Physics	48
3.1.3	Beam Control	49
3.1.4	Injection Sequence	50
3.1.5	LHC Run Schedules and Physics Performance	51
3.2	The CMS Detector	52
3.2.1	Overview of the Detector	52
3.2.2	The CMS coordinate system	53
3.2.3	The CMS magnet	54
3.2.4	Inner Trackers	55
	Pixel Tracker	56
	Strip Tracker	57
3.2.5	Electromagnetic Calorimetry	57
	Primary electromagnetic calorimeter	58
	Preshower	60
	Net Energy Resolution	61

3.2.6	Hadronic Calorimetry	61
	HCAL Barrel	62
	HCAL Endcaps	62
	Outer Calorimeter	63
	Forward Calorimeter	63
3.2.7	Muon Systems	63
	Drift Tubes	65
	Cathode Strip Chambers	65
	Resistive Plate Chambers	66
3.2.8	Trigger Systems	67
	L1 Trigger	68
	High Level Trigger	70
3.3	Physics Object Reconstruction	71
3.3.1	Tracks	72
3.3.2	Electrons and Photons	73
3.3.3	Muons	74
3.3.4	Jets	74
3.3.5	Missing transverse momentum p_T^{miss}	75
3.3.6	Particle Flow	76
4	Datasets and Triggers	77
4.1	Samples of Recorded Data	77
4.2	Simulated Samples	77
4.3	Triggers	81
4.4	Event Weights for Simulated Samples	83
5	Object Definitions and Event Selection	87
5.1	Object Definitions and Identification	87
5.1.1	Definitions of variables used for Particle ID	87
	ID variables for Jets	87
	ID variables for Photons	88
5.1.2	Jets	89
5.1.3	Photons	89
5.1.4	Missing transverse momentum p_T^{miss}	90
5.2	Key Analysis Variables	91
5.2.1	Scalar sum of transverse energy: S_T	91
5.2.2	Jet multiplicity: nJets	91
5.3	Event Selection Criteria	91
6	Background and Signal Models	96
6.1	S_T scaling	96
6.2	Background Model	98
6.2.1	Adaptive Gaussian Kernel estimation	98

	Optimizing ρ	99
	Expected number of events from S_T AGK template	100
6.2.2	Adjustments derived from MC	101
	Choosing a Functional Form for the Adjustment Term	102
	Best Fits from combined MC	103
	Finding adjustments from S_T AGK template	105
6.3	Signal Model	105
6.4	Full Model	105
6.5	Potential Signal Contamination	108
7	Uncertainties	116
7.1	Overview of Uncertainties	116
7.2	Uncertainties on Background Model	118
7.2.1	Statistical uncertainty on number of events in the normalization bin	118
7.2.2	Systematic uncertainty due to fluctuations in template shape	118
7.2.3	Systematic uncertainty due to ρ	120
7.2.4	Systematic uncertainty on MC-derived adjustments	120
7.2.5	Systematic uncertainties due to mismodeling	121
7.3	Uncertainties on Signal Model	129
7.3.1	Statistical uncertainty on weighted number of MC events	129
7.3.2	Systematic uncertainties due to jet energy correction, unclustered energy, and jet energy resolution effects	130
7.3.3	Systematic uncertainty due to prefiring weights	130
7.3.4	Systematic uncertainty due to photon MC scale factors	130
7.3.5	Uncertainty due HLT inefficiency	131
7.3.6	Systematic uncertainty due to luminosity	131
8	Statistical Analysis	132
8.1	General Overview	132
8.2	Procedure for Obtaining Upper Limits	133
8.2.1	Likelihood Function without nuisance parameters	133
8.2.2	Treatment of Nuisance Parameters	134
8.2.3	Test Statistic	136
8.2.4	Upper Limit Inference	137
9	Results	140
9.1	Overview	140
9.2	Expected Distributions and Expected Limits	141
9.2.1	Expected S_T shapes	141
9.2.2	Expected Limits	141
9.3	Fits to Observed Data	141
9.3.1	Observed S_T shapes	141

9.3.2	Nuisance Parameter Impacts	147
9.3.3	Validity of the Asymptotic Approximation	153
9.3.4	Analysis of Fits	153
9.4	Observed Limits and Interpretation	154
9.5	Concluding Remarks	157
9.5.1	Further Studies and Outlook	157
	Stealth SUSY with leptonic final states	158
	Stealth SUSY with displaced jets	159
Bibliography		160

List of Tables

1.1	Field Content of the SM	7
1.2	Transformations of the SM gauge fields under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$	8
1.3	Transformations of the SM leptonic fields under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$	8
1.4	Transformations of the SM quark fields under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$	9
1.5	Transformation of the SM Higgs field under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$	10
1.6	Chiral Superfields in the MSSM	29
1.7	Vector Superfields in the MSSM	29
4.1	Datasets and JSON masks used in this analysis.	78
4.2	MC datasets utilized in this Thesis.	79
4.3	Triggers used for data and background MC	83
5.1	Cut flow efficiencies, data	93
5.2	Cut flow efficiencies, Stealth MC	94
6.1	Various fits considered for the 2-to- n -jets transfer factor.	101
6.2	p -values from f -statistic for various fits.	102
6.3	Best-fit values of A and m used in background model.	103
7.1	Overview of uncertainties with approximate values.	117
9.1	Expected backgrounds, observations, sample signal yields.	145
9.2	Checking the validity of the asymptotic approximation.	153
9.3	Examining pre-fit vs post-fit error discrepancy for some nuisances. . .	154

List of Figures

1.1	Higgs mass correction terms for a fermionic Yukawa interaction. . . .	16
1.2	Running of the gauge couplings at two-loop level in the SM.	17
1.3	Correction term for the Higgs mass for a heavy complex scalar field. .	34
1.4	Running of the gauge couplings to leading order in the MSSM.	35
2.1	Some examples of SUSY models probed by collider experiments. . . .	38
2.2	Recent limits on gluino and squark production cross sections.	39
2.3	Key ingredients in models with Stealth SUSY.	41
2.4	Diagrams of the simplified models considered in this Thesis.	43
2.5	Exclusion limits on $\tilde{q}\tilde{q}$ production in Run 1 Stealth SUSY analysis. .	44
3.1	Schematic layout of the full accelerator complex.	47
3.2	Schematic layout of the LHC.	48
3.3	LHC integrated luminosity from 2011 to 2022.	51
3.4	A cutaway illustration of the CMS Detector.	52
3.5	Magnetic field of the CMS solenoid.	54
3.6	Illustration of the inner tracking system.	55
3.7	Illustration of ECAL and ES.	58
3.8	Illustration of HCAL.	61
3.9	Illustration of the muon systems.	64
3.10	The CMS L1 trigger workflow.	68
3.11	Trajectories of various particles through the detector.	72
4.1	HLT efficiencies for the diphoton trigger.	82
4.2	Distribution of prefiring probabilities.	84
4.3	Distribution of MC scale factors.	85
5.1	Search regions of the analysis, illustrating sidebands.	92
6.1	S_T scaling with QCD MC.	97
6.2	AGK Template at various values of ρ	99
6.3	3-fold NLL as a function of ρ	100
6.4	MC diphoton+ $(\gamma+\text{Jet})$ +multijet data and fits	104

6.5	Potential signal contamination for di-gluino events with 2 jets.	109
6.6	Potential signal contamination for di-squark events with 2 jets.	110
6.7	Potential signal contamination in the normalization range.	111
6.8	Relative signal correction, di-gluino events, $S_T < 1700$ GeV.	112
6.9	Relative signal correction, di-gluino events, $S_T > 1700$ GeV.	113
6.10	Relative signal correction, di-squark events, $S_T < 1700$ GeV.	114
6.11	Relative signal correction, di-squark events, $S_T > 1700$ GeV.	115
7.1	Toy MC datasets + templates for template shape systematic.	119
7.2	(shifted MC)/(nominal MC) as a function of S_T in the 4 jets bin.	122
7.3	(shifted MC)/(nominal MC) as a function of S_T in the 5 jets bin.	123
7.4	(shifted MC)/(nominal MC) as a function of S_T in the ≥ 6 jets bin.	123
7.5	S_T distributions, 1γ -control selection.	125
7.6	S_T distributions, 2γ -control selection.	126
7.7	2-to- n -jets adjustment in the 1γ -control selection.	127
7.8	2-to- n -jets adjustment in the 2γ -control selection.	128
7.9	Ratio of $(\frac{\text{data}}{\text{MC}} \text{ at } n \text{ jets})$ to $(\frac{\text{data}}{\text{MC}} \text{ at } 2 \text{ jets})$, 1γ -control selection	128
7.10	Ratio of $(\frac{\text{data}}{\text{MC}} \text{ at } n \text{ jets})$ to $(\frac{\text{data}}{\text{MC}} \text{ at } 2 \text{ jets})$, 2γ -control selection	129
9.1	Expected pre-fit background S_T distributions.	142
9.2	Expected upper limits on di-gluino production cross section.	143
9.3	Expected upper limits on di-squark production cross section.	144
9.4	Expected post-fit background S_T distributions and observed data.	146
9.5	Nuisance parameters: pulls and impacts (1/2).	147
9.6	Nuisance parameters: pulls and impacts (2/2).	148
9.7	Correlation matrix between nuisance parameters, b -only fit.	149
9.8	Correlation matrix between signal bins, b -only fit.	150
9.9	Correlation matrix between nuisance parameters, $s + b$ fit.	151
9.10	Correlation matrix between signal bins, $s + b$ fit.	152
9.11	Observed upper limits on di-gluino production cross section.	155
9.12	Observed upper limits on di-squark production cross section.	156
9.13	Stealth SUSY with leptonic final states.	158
9.14	Stealth SUSY interactions with vectorlike portal.	159

Organization and Notation

Organization

This Thesis is organized as follows.

Chapter 1 covers the theoretical background. It begins with a general reminder of the mathematical structure of the SM and describes the various issues that drive us to search for physics beyond the SM. Next we develop the mathematical structure of SUSY, and show how it can resolve a many of these issues.

In Chapter 2, we first discuss the general phenomenology of searches for SUSY at colliders and show that large parts of the “interesting” portion of SUSY parameter-space may have been ruled out already. Next we describe Stealth SUSY, a particular sub-class of SUSY models that, in fact, would not be covered by these searches.

Chapter 3 covers the “instruments” used to generate data for this analysis (although that word hardly does justice to the scale of the LHC and the CMS detector). At the end of this Chapter, we provide some details on basic physics object reconstruction and disambiguation, to give the reader a brief idea of how electronic signals are translated into a list of physics objects in a collision.

Chapter 4 covers the datasets (both simulated and recorded) used to obtain the signal and background yields, and covers the approach used to deal with potential inefficiencies.

Chapter 5 details the selections used to label collisions in order to define a subset optimized for reasonable signal-background discrimination. We also have other subsets with almost no potential expected signal, for use as control selections.

In Chapter 6, we describe the method used in this analysis to derive the background and signal models. This Chapter also discusses a particular subtle issue related to signal contamination that is somewhat of a unique feature of this analysis; the issue is resolved by the introduction of a correction term as described in the Chapter.

Chapter 7 describes the various sources of uncertainty on both the background and signal models, and describes how we estimate the magnitude of each.

Chapter 8, which can be read independently of the rest of this Thesis, documents how, in general, CMS analyses obtain and interpret limits on cross sections given a background model, a signal model, the observed data, and a set of systematic uncertainties, their magnitudes and their correlations.

Finally, Chapter 9 applies the methods of Chapter 8 to the Stealth SUSY search, and documents the results of that search.

Notation

Throughout this Thesis, unless indicated otherwise, we use natural units with $\hbar = c = 1$.

We use the “east cost”, mostly-plus metric: $\eta_{00} = -1$, $\eta_{11} = \eta_{22} = \eta_{33} = +1$.

Unless indicated otherwise, the Einstein convention is used: repeated indices are summed.

Greek indices $(\mu, \nu \dots)$ correspond to four Minkowski coordinates (t, x, y, z) . In general, Lorentz indices are raised or lowered with the metric: $\eta_{\mu\sigma}\eta^{\nu\tau}A^\mu_\nu = A^\tau_\sigma$.

Latin indices $(i, j \dots)$ stand for three spatial coordinates (x, y, z) .

The Pauli matrices are defined as:

$$\sigma_1 := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 := \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Undotted Latin indices $(a, b \dots)$ stand for the two spinor indices 1 and 2 of a left-handed spinor field (also called a “left-handed Weyl field”). Dotted Latin indices $(\dot{a}, \dot{b})$ stand for the two spinor indices $\dot{1}$ and $\dot{2}$ of a right-handed spinor field (or “right-handed Weyl field”). In general, left- or right-handed spinor indices are raised or lowered with the antisymmetric constants ε : $\chi_a = \varepsilon_{ab}\chi^b$, $\chi^a = \varepsilon^{ab}\chi_b$, $\chi^{\dagger a} = \varepsilon^{\dot{a}\dot{b}}\chi_{\dot{b}}^\dagger$. (Note that the antisymmetric nature of ε means that one has to be very careful about the order of the indices.) Here, we use these conventions:

$$\begin{aligned} \varepsilon^{12} = \varepsilon^{\dot{1}\dot{2}} = \varepsilon_{21} = \varepsilon_{\dot{2}\dot{1}} &= +1 \\ \varepsilon^{21} = \varepsilon^{\dot{2}\dot{1}} = \varepsilon_{12} = \varepsilon_{\dot{1}\dot{2}} &= -1 \end{aligned}$$

It is useful to define the following two constant “4-vectors” whose components are 2×2 matrices in spinor-space:

$$\sigma_{a\dot{a}}^\mu := (\mathbf{1}, \vec{\sigma}), \quad \bar{\sigma}^{\mu\dot{a}a} := (\mathbf{1}, -\vec{\sigma})$$

In order to avoid clutter, we generally omit spinor indices that are implicitly summed. In all such cases, we adopt the convention that a missing pair of undotted indices stands for the sum $\chi^a \psi_a$, while a missing pair of dotted indices stands for the sum $\chi_{\dot{a}} \psi^{\dot{a}}$. So, for example, the notations $\chi\psi$, $\chi^\dagger\psi^\dagger$, and $\psi^\dagger\bar{\sigma}^\mu\chi$ refer respectively to implied contractions $\chi^a\psi_a$, $\chi_{\dot{a}}^\dagger\psi^{\dagger\dot{a}}$, and $\psi_a^\dagger\bar{\sigma}^{\mu\dot{a}a}\chi_a$. With these conventions we always have $\psi\chi = \chi\psi$, $\psi^\dagger\chi^\dagger = \chi^\dagger\psi^\dagger$, and $(\chi\psi)^\dagger = \psi^\dagger\chi^\dagger$.

Right-handed Weyl fields are always written as the hermitian conjugates of a left-handed Weyl field (i.e. with a dagger) so that one can tell at a glance whether some suppressed indices are supposed to be dotted or undotted.

Most of these conventions follow the notation used in the textbook *Quantum Field Theory* by Mark Srednicki[1].

Chapter 1

Theoretical Background

The development of particle physics over the twentieth century can be characterized as a progressively more complete exploration of the key idea that fundamental models of Nature should be expressed in terms of as few parameters, and with as few assumptions, as possible. The Standard Model of Particle Physics (“SM”) is the result of this historical development. Section 1.1 presents a short informal introduction, and Section 1.2 contains a formal overview of the full mathematical structure of the SM. While it is exquisitely well-tested experimentally, there remain some issues, both philosophical and experimental, with the SM, as described in Section 1.3. The mathematical structure of a conjectured additional symmetry between quantum fields, called supersymmetry (“SUSY”), is described in Section 1.4, and Section 1.5 outlines why this particular conjecture is of special interest – and, specifically, how it is able to resolve some of the issues brought up in Section 1.3.

1.1 Introduction to the SM

Quantum field theory is the result of a successful attempt to reconcile two great paradigms in physics progressively developed over the first few decades of the twentieth century: special relativity, and quantum mechanics.

Special relativity is the notion that all observables should be invariant under a general Poincaré transformation – that is, under any combination of translations, rotations, or Lorentz boosts. More formally, any prediction made by a theory has to be invariant under the following transformation:

$$x^\mu \rightarrow x'^\mu = \Lambda^\mu_\nu x^\nu + a^\mu \tag{1.1}$$

where Λ (which can be some combination of a rotation and boost) is a 4×4 matrix that keeps the metric invariant:

$$\Lambda^\mu{}_\sigma \Lambda^\nu{}_\tau \eta_{\mu\nu} = \eta_{\sigma\tau} \quad (1.2)$$

and a (the spacetime translation) is an arbitrary 4-vector.

Quantum mechanics is the notion that the state of any physical system is describable by a vector $|\psi\rangle$ in Hilbert space. Physical observables are modeled as Hermitian operators in Hilbert space. The result of a measurement is always an eigenvalue λ_i of the operator, corresponding to the eigenvector $|\lambda_i\rangle$, and the Born rule states that the probability that the observable is measured to be λ_i is $|\langle\lambda_i|\psi\rangle|^2$. Symmetries (including, specifically, all Poincaré symmetries) are modeled as unitary operators in Hilbert space.

At first glance it might not seem obvious how to create a unified description of Nature that is both Poincaré-invariant and able to describe fundamental particles. However, one of the great triumphs of physics in the twentieth century has been to show that not only is it possible, there is in fact only one way to unify both ideas – and experiments confirm the extraordinary predictive power of this unique unification.

The starting point of quantum field theory is to note that the set of Poincaré transformations applied to Hilbert space, under usual composition rules for the unitary operators corresponding to those transformations, forms a ten-dimensional Lie group called the Poincaré group, whose Lie algebra is defined by the following commutation relations between the generators of translations (the momentum operators P_μ) and the generators of rotations and boosts (the operators $M_{\mu\nu}$):

$$\begin{aligned} [P_\mu, P_\nu] &= 0 \\ i [P_\mu, M_{\rho\sigma}] &= \eta_{\mu\rho} P_\sigma - \eta_{\mu\sigma} P_\rho \\ i [M_{\mu\nu}, M_{\rho\sigma}] &= \eta_{\mu\rho} M_{\nu\sigma} - \eta_{\mu\sigma} M_{\nu\rho} - \eta_{\nu\rho} M_{\mu\sigma} + \eta_{\nu\sigma} M_{\mu\rho} \end{aligned} \quad (1.3)$$

In this language, single-particle states are those whose transformations can be described by any one of the irreducible representations of the full Poincaré group. The precise transformations for each particle depend on the quantum numbers of the particle (the four-momentum, spin, helicity and so on).

The dynamical evolution of the state of a system is encoded in the Hamiltonian of the system (which is the generator of time translations). Causality is a very strong constraint on the structure of the Hamiltonian: the only known way to guarantee causality is to build the Hamiltonian from certain specific combinations of creation and annihilation operators that act on the vacuum to create and destroy free particle states. These special linear combinations are called “quantum fields”.

A Lagrangian formalism is used to derive the equations of motion for these quantum fields; a theory is specified completely by the parameters of the Lagrangian. Typically,

the observables of interest are quantities such as the interaction cross section for two initial state particles to scatter into a target final state, or the lifetime of an unstable particle. These are generally predicted by the theory in terms of the parameters of the Lagrangian using a perturbative expansion, and experiments impose progressively tighter constraints on these parameters.

1.2 The Mathematical Structure of the SM

1.2.1 SM Lagrangian

The Standard Model of particle physics is a quantum field theory that is invariant under the action of the gauge group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. All fields in the SM are listed in Table 1.1. Given this list of fields, one can construct a Lagrangian invariant under the gauge group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ as shown in the next few Sections.

The Lagrangian (before spontaneous symmetry breaking)

Ignoring the Higgs field for now, we can write the SM Lagrangian as follows:

$$\mathcal{L}_{\text{SM}}^0 = \mathcal{L}_{\text{gauge sector}}^0 + \mathcal{L}_{\text{lepton sector}}^0 + \mathcal{L}_{\text{quark sector}}^0 \quad (1.4)$$

In the next Sections, we shall describe each term.

Gauge sector

We begin with the gauge sector Lagrangian:

$$\mathcal{L}_{\text{gauge sector}}^0 = -\frac{1}{4}B^{\mu\nu}B_{\mu\nu} - \frac{1}{4}W^{I\mu\nu}W_{\mu\nu}^I - \frac{1}{4}F^{A\mu\nu}F_{\mu\nu}^A \quad (1.5)$$

where the meaning of the indices is in Table 1.1, and $B_{\mu\nu}$, $W_{\mu\nu}^I$, and $F_{\mu\nu}^A$ are built from covariant derivatives of the gauge fields:

$$\begin{aligned} B_{\mu\nu} &:= \partial_\mu B_\nu - \partial_\nu B_\mu \\ W_{\mu\nu}^I &:= \partial_\mu W_\nu^I - \partial_\nu W_\mu^I + g_2 \epsilon^{IJK} W_\mu^J W_\nu^K, \quad (I, J, K) \in (1, 2, 3) \\ F_{\mu\nu}^A &:= \partial_\mu F_\nu^A - \partial_\nu F_\mu^A + g_3 f^{ABC} F_\mu^B F_\nu^C, \quad (A, B, C) \in (1 \dots 8) \end{aligned} \quad (1.6)$$

where g_1 , g_2 , and g_3 are respectively the coupling constants for the $U(1)_Y$ hypercharge, $SU(2)_L$ isospin, and $SU(3)_C$ strong interactions. Here f^{ABC} are the structure constants of the $SU(3)$ Lie algebra, and we have used the fact that the structure constants for

Table 1.1: Field Content of the SM

Field	Notation	Description
$U(1)_Y$ gauge field	B	Gauge field mediating the $U(1)_Y$ electroweak hypercharge interaction. This is a vector field that carries a Lorentz index and is often expressed as B^μ .
$SU(2)_L$ gauge fields	W^I	Gauge fields mediating the $SU(2)_L$ electroweak isospin interactions. $I = 1, 2, 3$ correspond to the three $SU(2)$ generators. These are vector fields that carry a Lorentz index and are often expressed as $W^{I\mu}$.
$SU(3)_C$ gauge fields	F^A	Gauge fields mediating the $SU(3)_C$ strong interactions. $A = 1 \dots 8$ correspond to the eight $SU(3)$ generators. These are vector fields that carry a Lorentz index and are often expressed as $F^{A\mu}$.
Lepton $SU(2)_L$ -doublet	$\ell_I := \begin{pmatrix} \nu_I^L \\ e_I^L \end{pmatrix}$	$e_1^L, \nu_1^L, e_2^L, \nu_2^L, e_3^L$, and ν_3^L are left-handed Weyl fields and carry a spinor index (suppressed for clarity). These are respectively identified as the left-handed components of the electron, electron neutrino, muon, muon neutrino, tau, and tau neutrino fields.
Lepton $SU(2)_L$ -singlet	$e_I := e_I^R$	e_1^R, e_2^R , and e_3^R are left-handed Weyl fields and carry a spinor index (suppressed for clarity). $e_1^{R\dagger}, e_2^{R\dagger}$, and $e_3^{R\dagger}$ are respectively identified as the right-handed components of the electron, muon, and tau fields.
Quark $SU(2)_L$ -doublet	$q_I := \begin{pmatrix} u_I^L \\ d_I^L \end{pmatrix}$	$u_1^L, d_1^L, u_2^L, d_2^L, u_3^L, d_3^L$ are left-handed Weyl fields and carry a spinor index (suppressed for clarity). These are respectively identified as the left-handed components of the up, down, charm, strange, top, and bottom quark fields.
“Up”-type $SU(2)_L$ -singlet	$u_I := u_I^R$	u_1^R, u_2^R , and u_3^R are left-handed Weyl fields and carry a spinor index (suppressed for clarity). $u_1^{R\dagger}, u_2^{R\dagger}$, and $u_3^{R\dagger}$ are respectively identified as the right-handed components of the up, charm, and top quark fields.
“Down”-type $SU(2)_L$ -singlet	$d_I := d_I^R$	d_1^R, d_2^R , and d_3^R are left-handed Weyl fields and carry a spinor index (suppressed for clarity). $d_1^{R\dagger}, d_2^{R\dagger}$, and $d_3^{R\dagger}$ are respectively identified as the right-handed components of the down, strange, and bottom quark fields.
Higgs $SU(2)_L$ -doublet	$H := \begin{pmatrix} \chi \\ \phi \end{pmatrix}$	χ and ϕ are complex scalar fields.

Table 1.2: Transformations of the SM gauge fields under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$

Field	Transformation under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$	Additional Details
$B_\mu \rightarrow B'_\mu =$	$U_1(x)B_\mu U_1^\dagger(x) + \frac{i}{g_1}U_1(x)\partial_\mu U_1^\dagger(x),$ where $U_1(x) := \exp(-ig_1\Gamma_1(x)Y).$	Y is the hypercharge operator.
$W_\mu^I \rightarrow W_\mu^{I'} =$	$U_2(x)W_\mu^I U_2^\dagger(x) + \frac{i}{g_2}U_2(x)\partial_\mu U_2^\dagger(x),$ where $U_2(x) := \exp(-ig_2\Gamma_2^J(x)\tau^J).$	τ^J are the weak isospin operators.
$F_\mu^A \rightarrow F_\mu^{A'} =$	$U_3(x)F_\mu^A U_3^\dagger(x) + \frac{i}{g_3}U_3(x)\partial_\mu U_3^\dagger(x),$ where $U_3(x) := \exp(-ig_3\Gamma_3^B(x)T^B).$	T^B are the color charge operators.

the $SU(2)$ Lie algebra are the Levi-Civita symbol ϵ^{IJK} . Note that $U(1)_Y$ only has one generator and therefore no structure constants.

The transformations of the gauge fields under the SM gauge group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ are shown in Table 1.2. Under these transformations, the Lagrangian in Eq. 1.5 does not pick up any dependence on $\Gamma_1(x)$, $\Gamma_2^J(x)$, or $\Gamma_3^A(x)$, and we are free to choose any smooth functional form for these functions. This is the meaning of the statement that the Lagrangian in Eq. 1.5 is *invariant under the SM gauge group*.

Thinking of the gauge fields B , W , and F as general quantum fields, it is possible to imagine additional terms in the Lagrangian of the form $-\frac{1}{2}M_B^2 B^\mu B_\mu$, $-\frac{1}{2}M_{W^I}^2 W^{I\mu} W_\mu^I$, and $-\frac{1}{2}M_{F^A}^2 F^{A\mu} F_\mu^A$. These terms would respectively give a mass M_B , M_{W^I} , and M_{F^A} to the corresponding particles. However, all these terms would break gauge symmetry, since the Lagrangian would no longer be invariant under the transformations of Table 1.2. In general, in gauge-invariant theories, gauge particles are massless unless there are additional interactions in the theory.

Lepton sector

Table 1.3: Transformations of the SM leptonic fields under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$

Field	Transformation under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$
$\ell_I \rightarrow \ell'_I =$	$\exp(-ig_1(-\frac{1}{2})\Gamma_1(x)) \exp(-ig_2(+1)\Gamma_2^J(x)\frac{\sigma_I^J}{2})\ell_I$
$e_I \rightarrow e'_I =$	$\exp(-ig_1(+1)\Gamma_1(x))e_I$

The lepton sector Lagrangian has the following form:

$$\mathcal{L}_{\text{lepton sector}}^0 = i\ell_I^\dagger \bar{\sigma}^\mu D_\mu \ell_I + ie_I^\dagger \bar{\sigma}^\mu D_\mu e_I \quad (1.7)$$

where the action of the gauge-covariant derivative on the leptonic fields is defined as:

$$D_\mu \ell_I := \partial_\mu \ell_I - i \left\{ g_1 \times \left(-\frac{1}{2} \right) \right\} (B_\mu \ell_I) - i \{ g_2 \times (+1) \} \left(W_\mu^J \frac{\sigma_J}{2} \ell_I \right) \quad (1.8)$$

$$D_\mu e_I := \partial_\mu e_I - i \{ g_1 \times (+1) \} B_\mu e_I \quad (1.9)$$

Note that only the fields ℓ_I have a weak isospin charge. They are in the fundamental representation of $SU(2)_L$, and so the weak isospin generators are $\sigma_J/2$. The fields e_I are in the trivial representation of $SU(2)_L$. All these fields are in the trivial representation of $SU(3)_C$. Table 1.3 lists how the various lepton sector fields transform under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. Note that the Lagrangian in Eq. 1.7 is, by construction, gauge-invariant. It is possible to imagine that if the Lagrangian had a term such as $-m_I (e_I^L e_I^R + e_I^{R\dagger} e_I^{L\dagger})$, then the two-component spinors e_I^L and e_I^R could together

be identified as a four-component Dirac field $\begin{pmatrix} e_I^L \\ e_I^{R\dagger} \end{pmatrix}$ whose particles have mass m_I .

However, such terms in the Lagrangian would clearly not be invariant under the gauge transformations of Table 1.3, because the lepton doublets transform differently from singlets under $SU(2)_L$. Therefore, another mechanism needs to be at play in order for leptons to have a mass.

Quark sector

Table 1.4: Transformations of the SM quark fields under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$

Field	Transformation under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$
$q_I \rightarrow q'_I =$	$\exp \left(-ig_1 \left(+\frac{1}{6} \right) \Gamma_1(x) \right) \exp \left(-ig_2 (+1) \Gamma_2^J(x) \frac{\sigma_J}{2} \right) \exp \left(-ig_3 (+1) \Gamma_3^A(x) T^A \right) q_I$
$u_I \rightarrow u'_I =$	$\exp \left(-ig_1 \left(-\frac{2}{3} \right) \Gamma_1(x) \right) \exp \left(-ig_3 (+1) \Gamma_3^A(x) T^A \right) u_I$
$d_I \rightarrow d'_I =$	$\exp \left(-ig_1 \left(+\frac{1}{3} \right) \Gamma_1(x) \right) \exp \left(-ig_3 (+1) \Gamma_3^A(x) T^A \right) d_I$

The quark sector Lagrangian has the following form:

$$\mathcal{L}_{\text{quark sector}}^0 = i q_I^\dagger \bar{\sigma}^\mu D_\mu q_I + i u_I^\dagger \bar{\sigma}^\mu D_\mu u_I + i d_I^\dagger \bar{\sigma}^\mu D_\mu d_I \quad (1.10)$$

where the action of the gauge-covariant derivative on the quark fields is defined as:

$$D_\mu q_I := \partial_\mu \mathbb{1}_{q_I} - i \left\{ g_1 \times \left(+\frac{1}{6} \right) \right\} (B_\mu \mathbb{1}_{q_I}) - i \{ g_2 \times (+1) \} \left(W_\mu^J \frac{\sigma_J}{2} q_I \right) - i \{ g_3 \times (+1) \} (F_\mu^A T^A q_I) \quad (1.11)$$

$$D_\mu u_I := \partial_\mu u_I - i \left\{ g_1 \times \left(-\frac{2}{3} \right) \right\} B_\mu u_I - i \{ g_3 \times (+1) \} (F_\mu^A T^A u_I) \quad (1.12)$$

$$D_\mu d_I := \partial_\mu d_I - i \left\{ g_1 \times \left(+\frac{1}{3} \right) \right\} B_\mu d_I - i \{ g_3 \times (+1) \} (F_\mu^A T^A d_I) \quad (1.13)$$

Note that only the fields q_I have a weak isospin charge. They are in the fundamental representation of $SU(2)_L$, and so the weak isospin generators are $\sigma_J/2$. The fields u_I and d_I are in the trivial representation of $SU(2)_L$. Table 1.4 lists how the various quark sector fields transform under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. This Lagrangian is gauge-invariant by construction. Just as in the case of leptons, it is possible to imagine that if the Lagrangian had terms such as $-m_I (u_I^L u_I^R + u_I^{R\dagger} u_I^{L\dagger})$, then the two-component spinors u_I^L and u_I^R could together be identified a four-component Dirac field $\begin{pmatrix} u_I^L \\ u_I^{R\dagger} \end{pmatrix}$ whose particles have mass m_I (and the same would apply to the “down” fields d_I^L and d_I^R). However (also just as in the case of leptons), such terms in the Lagrangian would not be invariant under the gauge transformations of Table 1.4, because the quark doublets transform differently from the singlets under $SU(2)_L$. Therefore, another mechanism needs to be at play in order for quarks to have a mass.

Higgs Sector

Table 1.5: Transformation of the SM Higgs field under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$

Field	Transformation under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$
$\phi \rightarrow \phi' =$	$\exp \left(-ig_1 \left(-\frac{1}{2} \right) \Gamma_1(x) \right) \exp \left(-ig_2 (+1) \Gamma_2^J(x) \frac{\sigma_J}{2} \right) \phi$

In the previous Sections that introduced the Lagrangians for the gauge, lepton, and quark sectors, we saw that some additional complexity needs to be added in order to give masses to the various particles in the SM. In this Section, we introduce the Higgs field, which is an important component of the Higgs mechanism that provides these necessary mass terms. This Section describes the Higgs field; we describe the Higgs mechanism itself in more detail in Section 1.2.2.

We start by adding a new complex scalar field ϕ , called the Higgs field, which is an $SU(2)_L$ doublet. We give this doublet a so-called “Mexican hat” potential:

$$\mathcal{L}_{\text{Higgs}} = - (D^\mu \phi)^\dagger D_\mu \phi - \frac{\lambda}{4} \left(\phi^\dagger \phi - \frac{1}{2} v^2 \right)^2 \quad (1.14)$$

where $\lambda > 0$, and the action of the gauge-covariant derivative on the Higgs field is defined as:

$$D_\mu \phi := \partial_\mu \mathbb{1} \phi - i \left\{ g_1 \times \left(-\frac{1}{2} \right) \right\} (B_\mu \mathbb{1} \phi) - i \{ g_2 \times (+1) \} \left(W_\mu^J \frac{\sigma_J}{2} \phi \right) \quad (1.15)$$

The Higgs field is in the fundamental representation of $SU(2)_L$, so the weak isospin generators are $\sigma_J/2$. Table 1.5 shows how the Higgs field transforms under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. This Lagrangian is gauge-invariant by construction; uniquely for the Higgs field (as opposed to the quark or lepton fields), a direct mass term (implicit in the potential term in Eq. 1.14) does not break any gauge symmetry of the SM.

Yukawa couplings

Now that we have introduced the Higgs field and its properties under gauge transformations, we make an observation: while we cannot add direct mass terms to the Lagrangians in Eqs. 1.7 and 1.10, we can in fact add Yukawa coupling terms between the newly introduced Higgs field and the various lepton and quark fields. The nontrivial gauge transformation properties of the Higgs field makes it possible to add the following terms that are invariant under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$:

$$\mathcal{L}_{\text{lepton sector}}^{\text{Yukawa}} = -y_{IJ} (\phi \ell_I) e_J + \text{h.c.} \quad (1.16)$$

$$\mathcal{L}_{\text{quark sector}}^{\text{Yukawa}} = -y'_{IJ} (\phi q_I) d_J - y''_{IJ} (\phi^\dagger q_I) u_J + \text{h.c.} \quad (1.17)$$

where indices I and J run over the generations of leptons or quarks, and the parentheses represent an implicit contraction between the Higgs doublet and the lepton or quark doublet.

Full SM Lagrangian

We have now gathered all the terms necessary to write down the full Lagrangian of the SM, which is, by construction, Lorentz-invariant and gauge-invariant under $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$:

$$\boxed{\mathcal{L}_{\text{SM}}^{\text{full}} = \mathcal{L}_{\text{gauge sector}}^0 + \mathcal{L}_{\text{lepton sector}}^0 + \mathcal{L}_{\text{quark sector}}^0 + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{lepton sector}}^{\text{Yukawa}} + \mathcal{L}_{\text{quark sector}}^{\text{Yukawa}}} \quad (1.18)$$

where the terms $\mathcal{L}_{\text{gauge sector}}^0$, $\mathcal{L}_{\text{lepton sector}}^0$, $\mathcal{L}_{\text{quark sector}}^0$, $\mathcal{L}_{\text{Higgs}}$, $\mathcal{L}_{\text{lepton sector}}^{\text{Yukawa}}$, and $\mathcal{L}_{\text{quark sector}}^{\text{Yukawa}}$ can be read off Eqs. 1.5, 1.7, 1.10, 1.14, 1.16, and 1.17 respectively.

This is the complete Lagrangian of the SM. Gauge invariance is manifest in the form of Eq. 1.18, but this particular presentation comes at the price of obscuring a few key details – for one thing, it is not obvious from this form of the Lagrangian how particles other than the Higgs boson obtain their masses.

1.2.2 Spontaneous Symmetry Breaking

The acquisition of masses by the various SM particles (other than the Higgs boson) takes place through a general mechanism called “spontaneous symmetry breaking”, which crops up in many different areas of physics. The idea behind spontaneous symmetry breaking is that the vacuum or ground state of a theory need not be symmetric under the symmetries of the full Lagrangian. Since Nature is forced to pick a ground state that breaks the symmetry of the Lagrangian, this can introduce nontrivial dynamics, allowing effective terms into the Lagrangian that violate those symmetries.

In this case, the term responsible for symmetry breaking is the Higgs potential in Eq. 1.14, which gives the Higgs field a nonzero vacuum expectation value (“VEV”). We can always use a gauge transformation to rotate the VEV in $SU(2)_L$ -doublet-space and to change its phase. Here we adopt the convention that the VEV is entirely the first component of the Higgs doublet and is real:

$$\langle 0 | \phi(x) | 0 \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} v \\ 0 \end{pmatrix} \quad (1.19)$$

It can be seen that this particular gauge choice has broken gauge symmetry of the full Lagrangian corresponding to the subgroup $SU(2)_L \otimes U(1)_Y$ of the full SM gauge group. However, this gauge choice (called the “unitary gauge”) has not taken away all degrees of freedom: a key souvenir of $SU(2)_L \otimes U(1)_Y$ still persists, because any general $SU(2)_L$ gauge transformation along the 3-direction (but not the 1- or 2-directions) will only add a phase to the VEV; and this can be undone by a suitable choice of the $U(1)_Y$ gauge. In other words, there is a residual symmetry associated with the following operator:

$$Q^{\text{EM}} := Y + \tau^3 \quad (1.20)$$

where we have generalized the notation to apply to fields other than the Higgs field: Y is the hypercharge generator and τ^3 is the $SU(2)_L$ generator along the 3-direction – for SM fields, τ^3 is either the identity operator or $\sigma_3/2$ depending on the field. As hinted by the superscript, this is in fact the electromagnetic charge; we see that the

electromagnetic charge is conserved after symmetry breaking. We say that the full symmetry of the electroweak sector, $SU(2)_L \otimes U(1)_Y$, has been broken down into $U(1)_{EM}$.

In general, a complex doublet field has four real degrees of freedom. As we have seen, three of these degrees of freedom are taken away by the broken symmetry, but there is a residual choice of scale, which is carried by one real degree of freedom. The usual convention is to introduce this degree of freedom by expressing the Higgs field as an expansion about its VEV in terms of a single real field $H(x)$ as follows:

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} v + H(x) \\ 0 \end{pmatrix} \quad (1.21)$$

Making this substitution in Eq. 1.14 yields the following mass term in the Lagrangian in terms of a single real field $H(x)$:

$$\mathcal{L}_{\text{Higgs}}^{\text{mass terms}} = -\frac{1}{4}\lambda v^2 H^2 \quad (1.22)$$

(and the corresponding kinetic term is $-\frac{1}{2}\partial^\mu H \partial_\mu H$), which means that the Higgs boson picks up a mass equal to:

$$m_H := \sqrt{\frac{\lambda v^2}{2}} \quad (1.23)$$

We have seen that the nature of the Higgs potential forces us to fix the gauge to unitary gauge. Next we shall see what profound consequences this has for the SM.

Consequences of spontaneous symmetry breaking: gauge sector

Because the covariant derivative in Eq. 1.15 has terms that depend on the B and W gauge fields and are linear in ϕ , it is easy to see that a nonzero VEV for ϕ will lead to terms that are proportional to $v^2 BB$ and $v^2 WW$ when the Lagrangian in Eq. 1.14 is written out with ϕ expressed in unitary gauge as in Eq. 1.21. This is indeed what we see. First we give new names to certain linear combinations of the fields:

$$\begin{aligned} A_\mu &:= \sin \theta_W W_\mu^3 + \cos \theta_W B_\mu \\ Z_\mu &:= \cos \theta_W W_\mu^3 - \sin \theta_W B_\mu \\ W_\mu^\pm &:= \frac{1}{\sqrt{2}} (W_\mu^1 \mp i W_\mu^2) \end{aligned} \quad (1.24)$$

where $\theta_W := \arctan\left(\frac{g_1}{g_2}\right)$ is called the “Weinberg angle” (or “weak mixing angle”). With these replacements, the covariant derivative terms in Eq. 1.14 contribute the following mass terms to the full Lagrangian:

$$\mathcal{L}_{\text{gauge sector}}^{\text{mass terms}} = -M_W^2 W^{+\mu} W_\mu^- - \frac{1}{2} M_Z^2 Z^\mu Z_\mu \quad (1.25)$$

where $M_W := \frac{g_2 v}{2}$, and $M_Z := \frac{M_W}{\cos \theta_W}$. It is crucial to note that the field A does not pick up a mass term after spontaneous symmetry breaking. Indeed, this is not surprising – as we have seen, there is still a residual $U(1)_{EM}$ gauge symmetry after symmetry breaking, and a mass term for the corresponding field would violate this residual symmetry. Thus the A field is interpreted as the usual electromagnetic field and its particles are interpreted as photons (which, we emphasize again, are predicted to remain massless after spontaneous symmetry breaking). In general, the other electroweak gauge bosons observed in Nature are interpreted as the particles corresponding to these composite fields, and are all given special names: the particles corresponding to the Z and $W^\pm$ fields are called the Z (or Z^0) and $W^\pm$ bosons respectively.

Consequences of spontaneous symmetry breaking: lepton and quark sectors

We start with the lepton sector. In the SM, the Yukawa terms from Eq. 1.16 give masses to the leptons. This can be seen by examining Eq. 1.16. In general, it is possible by a biunitary transformation in generation-space to redefine all lepton fields such that the Yukawa matrix y_{IJ} is diagonal. The diagonal entry at index I then contributes the following term in unitary gauge (making the substitution from Eq. 1.21):

$$\mathcal{L}_{\text{lepton sector}}^{\text{mass terms}} = -\frac{1}{\sqrt{2}} y_I v (\ell_{I2} e_I) + \text{h.c.} \quad (1.26)$$

where the index 2 indicates that we need to take the second term of the doublet. Restoring the L and R identifiers from Table 1.1 to avoid confusion between the two different Weyl spinors, we can write Eq. 1.26 as follows:

$$\mathcal{L}_{\text{lepton sector}}^{\text{mass terms}} = -\frac{1}{\sqrt{2}} y_I v \left(e_I^L e_I^R + e_I^{R\dagger} e_I^{L\dagger} \right) \quad (1.27)$$

This is precisely the form of a Dirac mass term! We see that the Higgs mechanism induces leptons to pick up a mass $\frac{v y_I}{\sqrt{2}}$ from the Yukawa term in Eq. 1.16.

A very similar logic applies to the quark sector. But there arises a complication, because there are now two Yukawa terms instead of one; the biunitary transformation that diagonalizes y' is not necessarily the same as the one that diagonalizes y'' in

Eq. 1.17. It is still possible to diagonalize them separately, and indeed, these give rise to masses $\frac{vy'}{\sqrt{2}}$ and $\frac{vy''}{\sqrt{2}}$ for down-type and up-type quarks respectively, exactly as in the lepton sector. The additional complication, however, is that the gauge eigenstates are no longer equivalent to the mass eigenstates. This means there are “cross-currents” in the theory: a down-type quark can decay to an up-type quark from another generation (and vice versa) with the emission of a W with the appropriate charge. The degree of mixing is quantified by a 3×3 unitary matrix called the Cabibbo-Kobayashi-Maskawa (or CKM) matrix.

1.3 Issues with the SM

The SM has been spectacularly successful at describing the interactions between particles over a wide range of energy scales. Nonetheless, it has been known for some time that it is incomplete and not quite as satisfactory a model of natural phenomena as one might consider ideal. The next few paragraphs describe some issues with the SM.

1.3.1 Hierarchy Problem

From a philosophical point of view, one might ask why the observed masses of fermions are what they are. It might be expected that these masses should be tied to some meaningful physical scale, such as the Planck scale (10^{19} GeV).

In the SM, these masses are related to Yukawa coupling constants between the particle fields and the Higgs field, as described in Section 1.2.2. This explains why the masses of fundamental particles are fairly close to each other, since they are related to the Higgs scale by a dimensionless multiplicative constant¹, but this still does not explain why the observed Higgs mass of ~ 125 GeV is so much lower than the Planck mass. In other words, there are only two known fundamental mass scales in Nature: the Planck scale, and the electroweak or Higgs scale; and they are about 16 orders of magnitude different on a hypothetical hierarchy of mass scales. This surprising observation is called the “hierarchy problem”.

Naturally, in order to formalize this objection, one should first consider whether indeed there is any particular reason for the Higgs mass to be near the Planck mass (or, indeed, larger than its known value). As it turns out, in the SM, there actually is a basis for such an expectation. Consider Fig. 1.1, where f is any fermionic field with a Yukawa interaction with the Higgs field. In the SM, all fermionic fields have a term in the Lagrangian of the form $\lambda_f H \bar{f} f$, which results in the following correction to the

¹It is worth noting that the heaviest fermion (the top quark) is still about 10^5 times the mass of the lightest fermion (the electron). In that sense there is already a mini-hierarchy problem within the SM because the couplings vary over a few orders of magnitude.




Figure 1.1: Correction term for the Higgs mass for a Yukawa interaction with any fermionic field.

Higgs mass:

$$\Delta m_H^2 = -\frac{|\lambda_f|^2}{8\pi^2} \Lambda_{UV}^2 + \dots \quad (1.28)$$

where Λ_{UV} is an ultraviolet cutoff, set to some scale at which we expect physics beyond the general model to kick in. For example, the effects of gravity become important near the Planck scale m_{Pl} , so it is reasonable to set $\Lambda_{UV} \sim m_{Pl}$. This means that, without some new fields that add a correction term which precisely cancels the correction in Eq. 1.28, we would expect the Higgs mass to be around Λ_{UV} , regardless of the bare constants in the Higgs potential. However, for $\Lambda_{UV} \sim m_{Pl}$, the correction term is several orders of magnitude larger than the experimentally observed value of m_H^2 . In other words, if there is any physics at all beyond the Standard Model that involves an interaction with the Higgs field (a very reasonable expectation), it would require an extraordinary degree of collusion (up to 30 orders of magnitude!) between the bare constant in the Lagrangian and the correction term in order to produce the observed value of the Higgs mass. Since there is no explanation for this cancellation within the existing theory, the hierarchy problem is a strong indication that a new theory is needed.

It is interesting to note that while mass correction terms exist for all particles in the SM with nonzero mass, the correction terms for the gauge bosons and fermions only grow logarithmically in the cutoff scale Λ_{UV} , rather than quadratically as in the case of the Higgs field. In that sense, the fine-tuning is only manifestly an issue for the Higgs self-interaction. However, because the other particles do get their masses from the Higgs VEV, all fields in the SM are affected by this issue.




Figure 1.2: Running of the gauge couplings at two-loop level in the SM. Here, $\alpha_i := \frac{g_i^2}{4\pi}$. Credit: [2].

1.3.2 Running of the Gauge Couplings

Perturbative calculations are needed to translate the abstract theory to predictions of experimentally measurable quantities such as cross sections, particle lifetimes and so on. In these calculations, the experimentally observed values of the coupling constants are interpreted as a convolution of the “bare” coupling constants (such as the constants g_i implicit in Eq. 1.18) and various divergent terms. As long as there is a finite number of such divergent terms in the expansion, they can be absorbed into these bare constants. However, these divergent terms are generally expressed in terms of a renormalization scale, which is not a physical measurable. Typically, for collider experiments, it is convenient to choose the scale to be equal to the scale of the center-of-mass energy of the collision, which cancels out some terms in the perturbative expansion and makes calculations easier. The chosen renormalization scale is not a parameter of the model and no observation can depend on this scale; therefore, this also necessarily introduces an implicit scale dependence on what experimentalists interpret as the observed values of the coupling constants. This requirement that physical observables should not depend on the renormalization scale allows us to determine, based on the perturbative expansion, exactly how the observed values of the coupling constants are expected to evolve with the renormalization scale. This dependence, called the “running” of the gauge couplings, is shown for the SM in

Fig. 1.2 at the two-loop level.

It is entirely reasonable to believe that the SM by itself is not a complete description of Nature (it does not account for gravity, for example) and is instead simply an effective low-energy description of a deeper theory. The Higgs VEV breaks the full gauge symmetry of the SM and produces its diversity of fields; it is reasonable to expect that a similar mechanism should ensure that the SM itself should really be a low-energy description of a unified field whose full gauged structure is manifest at high energies. Practically, this means that the gauge couplings should meet each other at a so-called “grand unification” scale, which is the scale at which a unified model is needed to describe Nature. However, as shown in Fig. 1.2, the field content of the SM alone would not be consistent with such an expectation.

1.3.3 Neutrino Masses

In the SM, neutrinos are left-handed Weyl fields and there is no right-handed component (as opposed to the case for leptons or quarks). This means there is no allowed Yukawa coupling that is able to give a mass to the neutrinos after spontaneous symmetry breaking. However, it has been known for some time now [3, 4] that at least some neutrinos have a small, but nonzero, mass. While there are ways to account for nonzero neutrino masses within the SM, none of them are particularly satisfactory. For example, one could imagine introducing new left-handed Weyl fields that then become the right-handed components of the neutrino fields, which would then pick up mass via a Yukawa coupling with the Higgs field (in the same way that leptons and quarks obtain their masses). But this does not explain why neutrinos are so much lighter than the corresponding leptons, and relies on the existence of unobserved right-handed neutrinos. On the other hand, neutrinos could pick up mass by another mechanism altogether if they are their own antiparticles (which would be allowed because they have no EM charge) – in this case, they might be Majorana particles rather than Dirac particles. But this would lead to the same question regarding the unexpectedly small masses of the neutrinos, and in addition, would allow lepton number violation, which has not been observed so far.

1.3.4 Particle Nature of Dark Matter

So far in this discussion, we have neglected to mention gravity, because it plays no role in particle physics phenomenology at the scale of collider experiments. However, we do know that the general theory of relativity is a good description of Nature on large scales. Currently, we have a model of the universe on large scales that accounts for many observations in astrophysics and cosmology: galaxy rotation curves, cosmological expansion, gravitational lensing, the large-scale distribution of galaxies, and so on. This model is called the Λ CDM (“lambda-cold-dark-matter”) model, and it plays much the same role in astrophysics and cosmology that the SM does in particle physics.

The parameters of the model are fixed through cosmological observations (see, for example, recent results by the Planck collaboration [5]), and the best fit values for these parameters have strong implications. Of particular interest is the implication that roughly a quarter of the energy density of the universe is attributable to so-called “cold dark matter”. This is matter whose gravitational characteristics are exactly the same as ordinary matter. It is “cold” in the sense that the velocity of its constituent particles is much smaller than the speed of light at the redshift of matter-radiation equality (otherwise, galaxy formation would be disrupted at those redshifts), and it is mostly dissipationless in the sense that collisions do not carry away energy through SM particles such as photons (otherwise, dark matter would clump together rather than forming a diffuse halo around galaxies as observations would suggest).

However, given these known characteristics, the particle nature of dark matter remains a mystery. It is reasonable to expect that, just like ordinary matter and radiation, dark matter should also be describable by a quantum field theory, and if so, its excitations should be observable as particles; but there are no particles in the SM that would be a good candidate for cold dark matter. Neutrinos, while they do not interact much with the rest of the SM, are relativistic and would not be “cold”; all other fundamental SM particles are either unstable or have significant interactions with other SM particles.

1.4 Theoretical Underpinnings of SUSY

In the decades since the formulation of the SM as a complete theory in the early 1970s, various extensions have been proposed that attempt to address one or more of the issues with the SM described in Section 1.3. One particularly compelling conjecture is called “supersymmetry” (or SUSY), and in the next few Sections we shall develop the theory.

1.4.1 SUSY as an Extension of the SM

Coleman and Mandula showed [6] that, assuming some physically meaningful constraints such as nontriviality of the S -matrix, it is impossible to extend the Lie algebra of Eq. 1.3 in any theory in four spacetime dimensions in order to generate any new conserved charges that have a nontrivial geometric structure. In other words, Lorentz scalars such as the electromagnetic charge are allowed, but anything with a Lorentz index – besides the already known conserved quantities like the 4-momentum – are forbidden.

However, there is one possible loophole: conserved charges with spinorial indices would not violate the Coleman-Mandula theorem. Because this is the starting point of the SUSY conjecture, SUSY is particularly interesting from the theoretical perspective because it is one of the very few known ways to introduce nontrivial conservation laws into the laws of physics.

The generators of SUSY transformations are the operator Q_a (where a is a spinor index) and its Hermitian conjugate $Q_a^\dagger$, which extend the Poincaré algebra by adding the following commutation and anti-commutation relations to the Lie algebra in Eq. 1.3:

$$\begin{aligned}
\{Q_a, Q_a^\dagger\} &= -2\sigma_{aa}^\mu P_\mu \\
\{Q_a, Q_b\} &= 0 \\
\{Q_a^\dagger, Q_b^\dagger\} &= 0 \\
[Q_a, P_\mu] &= 0 \\
[Q_a^\dagger, P_\mu] &= 0 \\
[Q_a, M_{\mu\nu}] &= \frac{i}{4} (\sigma_\mu \bar{\sigma}_\nu - \sigma_\nu \bar{\sigma}_\mu)_a^b Q_b \\
[Q_a^\dagger, M_{\mu\nu}] &= -\frac{i}{4} Q_b^\dagger (\bar{\sigma}_\mu \sigma_\nu - \bar{\sigma}_\nu \sigma_\mu)^b_a
\end{aligned} \tag{1.29}$$

The commutators and anticommutators in Eq. 1.29, together with the usual commutators in Eq. 1.3, form a closed Lie algebra, referred to as the “superextension of the Poincaré group”. Because there is a direct connection between spacetime coordinates and the Poincaré group, it is not surprising that SUSY can be understood in a different context that has to do with additional spacetime coordinates. Usually, we consider quantum fields as being functions of the four spacetime variables; in SUSY, we instead work with “superfields” S , which, in addition to the four spacetime coordinates, also depend on complex two-component spinor coordinates (or Grassman variables) θ^a (and $\theta_a^\dagger$). The full set of coordinates $(x, \theta^a, \theta_a^\dagger)$ is called “superspace”. Given the properties of such spinorial variables (they have two independent components, and any term containing a product of more than two variables would be 0), the most general expansion of the superfield S reads as follows:

$$S(x, \theta, \theta^\dagger) = a + \theta\xi + \theta^\dagger\chi^\dagger + \theta\theta b + \theta^\dagger\theta^\dagger c + \theta^\dagger\bar{\sigma}^\mu\theta v_\mu + \theta^\dagger\theta^\dagger\theta\eta + \theta\theta\theta^\dagger\zeta^\dagger + \theta\theta\theta^\dagger\theta^\dagger d \tag{1.30}$$

In the most general case, a , b , c , and d are complex functions of position that can be thought of as bosonic complex fields; they each have 2 degrees of freedom. The complex-vector-valued function of position v_μ can be thought of as a bosonic complex vector field; it has 8 degrees of freedom. The quantities ξ , $\chi^\dagger$, η , and $\zeta^\dagger$ are complex-spinor-valued functions of position and can be thought of as fermionic fields; they each have 4 degrees of freedom. On the other hand, if the superfield S is constrained to be real-valued, then a , b , c , d , are all real-valued with one degree of freedom each;

v_μ is real-valued with four degrees of freedom; and we find that $\xi = \chi$ and $\eta = \zeta$, eliminating two fermionic degrees of freedom. In both cases, we have exactly the same number of bosonic degrees of freedom and fermionic degrees of freedom, which is an important observation.

This motivates a SUSY construction called a “supermultiplet”. There are two types of supermultiplets, for which we shall derive the Lagrangian in the next few Sections:

- Chiral supermultiplets contain one fermion with two helicity states, and one complex scalar boson with two components.
- Gauge supermultiplets contain one vector boson with two spin states (in the unbroken theory, the vector boson is massless), and one fermion with two helicity states.

The ingredients used to build these supermultiplets are chiral superfields and vector superfields.

Chiral Superfields

Just as we used gauge-covariant derivatives to build the SM Lagrangian, so also we can usefully define the following “chiral-covariant” derivatives in SUSY:

$$\begin{aligned} D_a &= \frac{\partial}{\partial \theta^a} - i (\sigma^\mu \theta^\dagger)_a \partial_\mu \\ D^a &= -\frac{\partial}{\partial \theta_a} + i (\theta^\dagger \bar{\sigma}^\mu)^a \partial_\mu \end{aligned} \tag{1.31}$$

and

$$\begin{aligned} \bar{D}^{\dot{a}} &= \frac{\partial}{\partial \theta_a^\dagger} - i (\bar{\sigma}^\mu \theta)^{\dot{a}} \partial_\mu \\ \bar{D}_{\dot{a}} &= -\frac{\partial}{\partial \theta^{\dagger \dot{a}}} + i (\theta \sigma^\mu)_{\dot{a}} \partial_\mu \end{aligned} \tag{1.32}$$

Of particular interest are those superfields Φ that belong to the kernels of the operators defined by Eqs. 1.31 and 1.32. In particular, we define a “left-chiral” superfield as a field that satisfies the equation:

$$\bar{D}_{\dot{a}} \Phi = 0 \tag{1.33}$$

One can show that with these conventions, its complex conjugate Φ^* obeys $D_a \Phi^* = 0$ and is called a “right-chiral” superfield. One can also show that the general solution to Eq. 1.33 can be expressed as follows:

$$\begin{aligned}\Phi(x, \theta, \theta^\dagger) = & \phi(x) + i\theta^\dagger \bar{\sigma}^\mu \theta \partial_\mu \phi(x) + \frac{1}{4} \theta \theta \theta^\dagger \theta^\dagger \partial_\mu \partial^\mu \phi(x) \\ & + \sqrt{2} \theta \psi(x) - \frac{i}{\sqrt{2}} \theta \theta \theta^\dagger \bar{\sigma}^\mu \partial_\mu \psi(x) + \theta \theta F(x)\end{aligned}\tag{1.34}$$

The superfield defined by Eq. 1.34 has four degrees of freedom: a complex scalar field ϕ , a Weyl spinor ψ (which has two components), and an auxiliary field F .

Vector Superfields

Vector superfields are defined by simply requiring the superfield S to be real-valued, in which case Eq. 1.30 can be written as:

$$S(x, \theta, \theta^\dagger) = a + \theta \chi + \theta^\dagger \chi^\dagger + \theta \theta b + \theta^\dagger \theta^\dagger c + \theta^\dagger \bar{\sigma}^\mu \theta v_\mu + \theta^\dagger \theta^\dagger \theta \eta + \theta \theta \theta^\dagger \eta^\dagger + \theta \theta \theta^\dagger \theta^\dagger d \tag{1.35}$$

where a, b, c, d are real-valued functions, v_μ is a real-vector-valued function, and χ and η are fermionic fields. As discussed earlier, there are eight bosonic and eight fermionic degrees of freedom here; however, we note here that there is a gauge redundancy implicit in Eq. 1.35: it is possible to show that by a choice of gauge, we can eliminate four bosonic and four fermionic degrees of freedom. In this gauge (called the Wess-Zumino gauge), a general vector superfield has the following simpler form:

$$V_{\text{WZ}} = \theta^\dagger \bar{\sigma}^\mu \theta A_\mu + \theta^\dagger \theta^\dagger \theta \lambda + \theta \theta \theta^\dagger \lambda^\dagger + \frac{1}{2} \theta \theta \theta^\dagger \theta^\dagger D \tag{1.36}$$

1.4.2 SUSY Lagrangians

It is easy to see that products of an arbitrary number of superfields in the form of Eq. 1.30 are also superfields in the form of Eq. 1.30. This means we have already developed some of the language necessary to build SUSY Lagrangians corresponding to supermultiplets.

In general, the translation between SUSY theories and ordinary quantum field theory is performed by integrating out the fermionic degrees of freedom from a general superfield, leaving behind a Lagrangian density in the four usual spacetime coordinates. Our goal is to do this in a way that is invariant under SUSY transformations. The following constructions prove useful:

F -term

For any chiral superfield Φ , we can define its F -term as follows:

$$\Phi_F := \int d^2\theta \Phi(\theta^\dagger = 0) = \int d^2\theta d^2\theta^\dagger \delta^2(\theta^\dagger) \Phi(x, \theta, \theta^\dagger) = F \quad (1.37)$$

where F is the degree of freedom associated with the $\theta\theta$ term in Eq. 1.34. To keep the action real, such a term needs to be added to its complex conjugate before adding it to the Lagrangian. It can be shown that such a contribution to the action is invariant under superspace transformations generated by Q and $Q^\dagger$ in Eq. 1.29.

D -term

For any vector superfield V , we can define its D -term as follows:

$$V_D := \int d^2\theta d^2\theta^\dagger V(x, \theta, \theta^\dagger) = d \quad (1.38)$$

where d is the degree of freedom associated with the $\theta\theta\theta^\dagger\theta^\dagger$ term in Eq. 1.35. Such a term can be added directly to the Lagrangian. Just as for the F -term, it can be shown that such a contribution to the action is also invariant under general SUSY transformations in superspace.

SUSY Lagrangian with no gauging

In order to build a SUSY Lagrangian, we can either take the D -term of a vector superfield or the F -term of a chiral superfield plus its complex conjugate.

We can now start examining the form of a theory built from these constructions. First, suppose there are multiple left-chiral superfields Φ_i . The product $\Phi_i^* \Phi_j$ is itself a superfield that might be expected to describe interactions between the composite fields of each of Φ_i^* and Φ_j . Consider just the terms with $i = j$ first; clearly, $\Phi_i^* \Phi_i$ is a vector superfield, and so it would be expected to contribute a D -term to the full Lagrangian, which can be read off as the coefficient of the $\theta\theta\theta^\dagger\theta^\dagger$ term:

$$(\Phi_i^* \Phi_i)_D = -\partial^\mu \phi_i^* \partial_\mu \phi_i + i\psi_i^\dagger \bar{\sigma}^\mu \partial_\mu \psi_i + F_i^* F_i + \text{total spacetime derivatives} \quad (1.39)$$

where the total spacetime derivatives can be ignored since they will not feature in the action after integration. This form is somewhat encouraging because it already features the basic necessities of a supermultiplet: kinetic terms in scalar and spinor fields.

The mass and interaction terms can be obtained with the realization that the product of two chiral superfields is a chiral superfield, which means we need to add F -terms for these products. More generally, any holomorphic function $\mathcal{W}$ of a product of superfields (which, in this context, means a function that depends on Φ but not Φ^*) could feature in the full Lagrangian, which can be written as:

$$\mathcal{L}(x) = (\Phi_i^* \Phi_i)_D + (\mathcal{W}(\Phi))_F + \text{c.c.}) \quad (1.40)$$

If the full theory is to be renormalizable, however, it can be shown that there are some constraints on $\mathcal{W}$: it can only have the following terms²:

$$\mathcal{W}(\Phi) = \frac{1}{2} M_{ij} \Phi_i \Phi_j + \frac{1}{6} y_{ijk} \Phi_i \Phi_j \Phi_k \quad (1.41)$$

Gauging the SUSY Lagrangian

The next step is to generalize this formalism by gauging the superfield. We start with a single gauge group; we have to ensure invariance under so-called “supergauge” transformations, defined as:

$$\begin{aligned} \Phi_i &\rightarrow \Phi'_i = (\exp(2ig\Omega^I T^I))_{ij} \Phi_j \\ \Phi_i^* &\rightarrow \Phi'^*_i = \Phi_j^* (\exp(-2ig\Omega^I T^I))_{ji} \end{aligned} \quad (1.42)$$

where g is the coupling constant, Ω^I are chiral superfields, and T^I are the generators of the supergauge transformation with a Lie algebra given by $[T^I, T^J] = if^{IJK} T^K$ and normalized to $2\text{Tr}(T^I T^J) = \delta^{IJ}$, as in any general non-Abelian gauge theory. In order to keep the full analog of the kinetic term supergauge-invariant in the Lagrangian in Eq. 1.40, it has to be modified by introducing new vector superfields V^I :

$$\mathcal{L}_{\text{kin}}(x) = \left(\Phi_i^* \exp(2gT^I V^I)_{ij} \Phi_j \right)_D \quad (1.43)$$

where the superfields V^I have supergauge transformations that obey the relation:

$$\exp(2gT^I V^I) \rightarrow \exp(2gT^I V'^I) = \exp(2igT^I \Omega^{I\dagger}) \exp(2gT^I V^I) \exp(-2igT^I \Omega^I) \quad (1.44)$$

²Technically, a term $\ell_i \Phi_i$ would also be allowed by renormalizability, but it would imply the existence of a gauge singlet, which is not present at least in those models of SUSY that are based on minimal extensions to the SM.

It is possible to stick to this convention in the Wess-Zumino gauge, which simplifies the expressions somewhat. In the Wess-Zumino gauge, the superfields V^I have the expression:

$$V_{\text{WZ}}^I = \theta^\dagger \bar{\sigma}^\mu \theta A_\mu^I + \theta^\dagger \theta^\dagger \theta \lambda^I + \theta \theta \theta^\dagger \lambda^{I\dagger} + \frac{1}{2} \theta \theta \theta^\dagger \theta^\dagger D^I \quad (1.45)$$

If ∇_μ denotes the action of the supergauge-covariant derivatives³, then the kinetic term of Eq. 1.43 (expressing Φ_i in terms of its component fields ϕ_i , ψ_i , and F_i , and V^I in terms of Eq. 1.45) can be expressed in a manifestly supergauge-invariant manner as follows:

$$\begin{aligned} \mathcal{L}_{\Phi\text{-fields+interactions}}(x) &= \left(\Phi_i^* \exp(2gT^I V^I)_{ij} \Phi_j \right)_D \\ &= -\nabla_\mu \phi_i^* \nabla^\mu \phi_i + i\psi_i^\dagger \bar{\sigma}^\mu \nabla_\mu \psi_i - \sqrt{2}g(\phi_i^* T^I \psi_i) \lambda^I - \sqrt{2}g\lambda^{I\dagger}(\psi_i^\dagger T^I \phi_i) \\ &\quad + g(\phi_i^* T^I \phi_i) D^I + F_i^* F_i \end{aligned} \quad (1.46)$$

However, this is not yet the complete set of all kinetic terms: we need to introduce kinetic terms for the vector supermultiplets, and this is done as in general non-Abelian gauge theory by defining the following chiral superfield⁴:

$$W_a = -\frac{1}{4} \overline{D}\overline{D} \left(\exp(-2gT^I V^I) D_a \exp(2gT^I V^I) \right) \quad (1.47)$$

which, under a supergauge transformation, transforms as:

$$W_a \rightarrow W'_a = \exp(2ig\Omega^I T^I) W_a \exp(-2ig\Omega^I T^I) \quad (1.48)$$

One can recover components W_a^I of W_a in the adjoint representation (where the components are defined by $W_a := 2gT^I W_a^I$), and express each chiral field in Wess-Zumino gauge, leading to:

$$W_a^I = \lambda_a^I + \theta_a D^I + \frac{i}{2} (\sigma^\mu \bar{\sigma}^\nu \theta)_a F_{\mu\nu}^I + i\theta\theta (\sigma^\mu \nabla_\mu \lambda^{I\dagger})_a \quad (1.49)$$

where F is the usual non-Abelian gauge field:

³We use this notation instead of the plain D_μ in Section 1.2 to distinguish it from the D term in a vector superfield as well as the gauge-covariant differential operators D_a and $\overline{D}_a$ in Eqs. 1.31 and 1.32.

⁴It can be shown that $\overline{D}\overline{D}S$ is a chiral superfield for any arbitrary superfield S .

$$F_{\mu\nu}^I := \partial_\mu A_\nu^I - \partial_\nu A_\mu^I - gf^{IJK} A_\mu^J A_\nu^K \quad (1.50)$$

Finally, one can obtain the remaining kinetic terms that are also supergauge-invariant by extracting the relevant F -term:

$$\begin{aligned} \mathcal{L}_{V\text{-fields+interactions}}(x) &= \frac{1}{2} (W_a^I W_a^I)_F \\ &= \frac{1}{2} D^I D^I + i\lambda^I \sigma^\mu \nabla_\mu \lambda^{\dagger I} - \frac{1}{4} F^{I\mu\nu} F_{\mu\nu}^I + \frac{i}{8} \epsilon^{\mu\nu\rho\sigma} F_{\mu\nu}^I F_{\rho\sigma}^I \end{aligned} \quad (1.51)$$

There is one more simplification that needs to be made before we write down the full Lagrangian. It can be seen from the Lagrangians in Eqs. 1.46 and 1.51 that the fields F_i and D^I do not have kinetic terms. The equations of motion for both fields are purely algebraic relations between these fields and others in the full theory; carrying out these replacements yields some important insights. In particular, with the form of $\mathcal{W}$ fixed by Eq. 1.41, the equations of motion imply the following:

$$\begin{aligned} F_i &= -\mathcal{W}_i^* := -M_{ij}^* \phi_j^* - \frac{1}{2} y_{ijk}^* \phi_j^* \phi_k^* \\ F_i^* &= -\mathcal{W}_i := -M_{ij} \phi_j - \frac{1}{2} y_{ijk} \phi_j \phi_k \\ D^I &= -g (\phi_i^* T^I \phi_i) \end{aligned} \quad (1.52)$$

With these substitutions, the full Lagrangian takes the following form:

$$\begin{aligned} \mathcal{L}_{\text{full}} &= -\nabla_\mu \phi_i^* \nabla^\mu \phi_i + i\psi_i^\dagger \bar{\sigma}^\mu \nabla_\mu \psi_i \\ &\quad + i\lambda^I \sigma^\mu \nabla_\mu \lambda^{\dagger I} - \frac{1}{4} F^{I\mu\nu} F_{\mu\nu}^I + \frac{i}{8} \epsilon^{\mu\nu\rho\sigma} F_{\mu\nu}^I F_{\rho\sigma}^I \\ &\quad - \sqrt{2}g(\phi_i^* T^I \psi_i) \lambda^I - \sqrt{2}g\lambda^{\dagger I} (\psi_i^\dagger T^I \phi_i) \\ &\quad - \frac{g^2}{2} (\phi_i^* T^I \phi_i)^2 \\ &\quad - (M_{ij} \phi_j + \dots) (M_{ik}^* \phi_k^* + \dots) \\ &\quad - \frac{1}{2} \left(M_{ij} \psi_i \psi_j + M_{ij}^* \psi_i^\dagger \psi_j^\dagger \right) + \dots \end{aligned} \quad (1.53)$$

where the last term above comes from the F -term of the product $\Phi_i \Phi_j$ and we have suppressed the higher order interaction terms for brevity.

Lagrangian with full gauge group

Eq. 1.53 is the full Lagrangian for a gauged theory with with one vector superfield and one coupling constant g . If, instead, the full gauge group of Nature happens to be a direct product of gauge subgroups (such as the SM $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$), that can be accommodated without additional subtlety into the Lagrangian in Eq. 1.53. Now, Eq. 1.42 needs multiple terms in front of Φ_j , one for each gauge group; indexing the gauge group by A (so that the coupling constant for the gauge group is g_A , its generators are given by T_A^I , etc.), the full Lagrangian can be written as:

$$\begin{aligned}
\mathcal{L}_{\text{full}}^{\text{SUSY}} = & -\nabla_\mu \phi_i^* \nabla^\mu \phi_i + i\psi_i^\dagger \bar{\sigma}^\mu \nabla_\mu \psi_i \\
& + i\lambda_A^I \sigma^\mu \nabla_\mu \lambda_A^{\dagger I} - \frac{1}{4} F_A^{I\mu\nu} F_{A\mu\nu}^I + \frac{i}{8} \epsilon^{\mu\nu\rho\sigma} F_{A\mu\nu}^I F_{A\rho\sigma}^I \\
& - \sqrt{2}g_A (\phi_i^* T_A^I \psi_i) \lambda_A^I - \sqrt{2}g_A \lambda_A^{\dagger I} (\psi_i^\dagger T_A^I \phi_i) \\
& - \frac{g_A^2}{2} (\phi_i^* T_A^I \phi_i)^2 \\
& - (M_{ij} \phi_j + \dots) (M_{ik}^* \phi_k^* + \dots) \\
& - \frac{1}{2} \left(M_{ij} \psi_i \psi_j + M_{ij}^* \psi_i^\dagger \psi_j^\dagger \right) + \dots
\end{aligned} \tag{1.54}$$

where, as earlier, we have suppressed the higher order terms (that depend on y_{ijk}) for brevity.

1.4.3 Mass degeneracy of superpartners in SUSY

An important result can be derived by examining the full Lagrangian in Eq. 1.54. The following terms quadratic in ϕ are responsible for giving a mass to the ϕ scalars:

$$\mathcal{L}_{\phi\text{-mass}} = -M_{ik}^* M_{kj} \phi_i^* \phi_j \tag{1.55}$$

The mass spectrum of the scalars of this theory is obtained by diagonalizing the positive semidefinite matrix $M^* M$, whose eigenvalues are the squares of the masses in the spectrum.

The fermionic mass terms are:

$$\mathcal{L}_{\psi\text{-mass}} = -\frac{1}{2} \left(M_{ij} \psi_i \psi_j + M_{ij}^* \psi_i^\dagger \psi_j^\dagger \right) \tag{1.56}$$

and lead to the following equations of motion:

$$\begin{aligned} i\bar{\sigma}^\mu \partial_\mu \psi_i - M_{ij}^* \psi_j^\dagger &= 0 \\ i\sigma^\mu \partial_\mu \psi_i^\dagger - M_{ij}^* \psi_j &= 0 \end{aligned} \tag{1.57}$$

which can be combined (by applying a differential operator on to any one equation and substituting the other):

$$\partial_\mu \partial^\mu \psi_i + M_{ij}^* M_{jk} \psi_k = 0 \tag{1.58}$$

This means that the mass spectrum of the fermions of this theory is also obtained by diagonalizing the same positive semidefinite matrix $M^* M$, whose eigenvalues are the squares of the masses of the fermions. Given that the number of fermionic and scalar fields in any SUSY theory is the same by construction, this means that the mass spectrum of scalars and fermions is exactly identical – and this is simply a natural consequence of the structure of the theory. If SUSY is an exact symmetry, then there is an exact degeneracy between the masses of all scalars and fermions in the theory!

1.4.4 The Minimal Supersymmetric Standard Model (MSSM)

In order for a SUSY model to match the success of the SM, it is desirable to use its structure as the starting point for the theory. The Minimal Supersymmetric Standard Model (MSSM) extends the SM by promoting each of its fields to a superfield, thus giving the corresponding particle a superpartner. SM fermions (lepton doublets and singlets, quark doublets and singlets) are taken to be the fermionic components of a chiral superfield; these chiral supermultiplets contain superpartners of the fermions, which are all spin 0 scalars. SM gauge fields (B for $U(1)_Y$, W^I for $SU(2)_L$, and F^A for $SU(3)_C$) are all taken to be the vector components of vector superfields; these gauge supermultiplets contain superpartners of the vectors, which are all spin $1/2$ fermions. In the MSSM (unlike in the SM), two Higgs scalars are required to break the $SU(2)_L \otimes U(1)_Y$ symmetry for a reason we shall discover in a few paragraphs; these are taken to be the scalar components of a chiral superfield, and the corresponding chiral supermultiplets contain two spin $1/2$ fermions. They are labeled “u” and “d” because the VEVs of these scalars are responsible for the masses of the “up”-type and “down”-type quarks respectively.

Tables 1.6 and 1.7 contain the complete list of MSSM chiral and vector superfields respectively. There is a notation that has become conventional in discussions of SUSY: the scalar partners of the fermions have an “s-” prefix and are called “sfermions” (so, for example, the superpartner of the electron is the selectron, the superpartner of the top quark is the stop, the superpartners of quarks in general are squarks, etc.); the fermionic partners of the gauge particles have an “-ino” suffix and are called “gauginos”

Table 1.6: Chiral Superfields in the MSSM

Superfield	Notation (super-field)	Scalar	Spinor	Description
Lepton $SU(2)_L$ - doublet	L_I	$\tilde{\ell}_I := \begin{pmatrix} \tilde{\nu}_I^L \\ \tilde{e}_I^L \end{pmatrix}$	$\ell_I := \begin{pmatrix} \nu_I^L \\ e_I^L \end{pmatrix}$	I runs over three lepton generations. The spinor particles are the leptons from Table 1.1, while the scalar particles are called “sleptons”.
Lepton $SU(2)_L$ - singlet	E_I	$\tilde{e}_I := \tilde{e}_I^R$	$e_I := e_I^R$	
Quark $SU(2)_L$ - doublet	Q_I	$\tilde{q}_I := \begin{pmatrix} \tilde{u}_I^L \\ \tilde{d}_I^L \end{pmatrix}$	$q_I := \begin{pmatrix} u_I^L \\ d_I^L \end{pmatrix}$	
“Up”-type $SU(2)_L$ - singlet	U_I	$\tilde{u}_I := \tilde{u}_I^R$	$u_I := u_I^R$	I runs over three quark generations. The spinor particles are the quarks from Table 1.1, while the scalar particles are called “squarks”.
“Down”-type $SU(2)_L$ - singlet	D_I	$\tilde{d}_I := \tilde{d}_I^R$	$d_I := d_I^R$	
“Up”-type Higgs	H_u	$h_u := \begin{pmatrix} H_u^+ \\ H_u^0 \end{pmatrix}$	$\tilde{h}_u := \begin{pmatrix} \tilde{H}_u^+ \\ \tilde{H}_u^0 \end{pmatrix}$	The MSSM requires two Higgs scalars as described in the text. The spinor fields are called “Higgsinos”.
“Down”-type Higgs	H_d	$h_d := \begin{pmatrix} H_d^0 \\ H_d^- \end{pmatrix}$	$\tilde{h}_d := \begin{pmatrix} \tilde{H}_d^0 \\ \tilde{H}_d^- \end{pmatrix}$	

Table 1.7: Vector Superfields in the MSSM. Note: the overlines on the superfield symbols are part of the name and do not imply the usual conjugation.

Superfield	Notation (super-field)	Vector	Spinor	Description
$U(1)_Y$ gauge field	$\overline{B}$	B	$\tilde{B}$	The spinor partner of the B field is called a “bino”.
$SU(2)_L$ gauge fields	$\overline{W}^I$	W^I	$\widetilde{W}^I$	I runs over the three $SU(2)$ generators. The spinor partners of the W fields are called “winos”.
$SU(3)_C$ gauge fields	$\overline{F}^A$	F^A	$\tilde{F}^A$	A runs over the eight $SU(3)$ generators. The spinor partners of the gluons are called “gluinos”.

(so, for example, the superpartner of the photon is the photino, the superpartner of the gluon is the gluino, etc.). The symbols for SUSY partner fields are usually the symbols of the corresponding SM fields with a tilde on top.

The Lagrangian in Eq. 1.54 demonstrates that the structure of a given SUSY model is characterized entirely by a limited number of free parameters: in the discussion leading up to Eq. 1.54, these interactions were encoded in $\mathcal{W}$ – which is also called the “superpotential” – and, by extension, in the parameters M and y . Of course, in the full theory, there are multiple superfields that all interact with each other. The superpotential of the MSSM, which defines the interaction terms in the full theory, is expressed as follows:

$$\mathcal{W}_{\text{MSSM}} = U_I Y_{IJ}^U Q_J H_u - D_I Y_{IJ}^D Q_J H_d - E_I Y_{IJ}^E L_J H_d + \mu H_u H_d \quad (1.59)$$

where the indices I, J run over three generations each for the quark/squark and lepton/slepton superfields. The constant coefficients Y_{IJ}^U , Y_{IJ}^D , and Y_{IJ}^E are Yukawa terms that are 3×3 matrices that, after H_u and H_d get a VEV, yield masses and mixing angles between generations of quarks and leptons. The last term, called the “ μ -term”, deserves some further explanation. In the SM, with only one Higgs field we could construct a Lagrangian that obeyed all the symmetries of the system. But consider a term like $H_u^* H_u$ or $H_d^* H_d$ (which would each be a direct analog of the SM Higgs). While it would be gauge-invariant, there is no straightforward way to make it invariant under SUSY: recall that we can only take the F -term of a holomorphic function of a product of left-chiral superfields (not their complex conjugates, which are right-chiral superfields) in building the superpotential, as discussed in Section 1.4.2. This is the reason the MSSM needs two Higgs fields with hypercharge assignments $\pm \frac{1}{2}$. This means that after electroweak symmetry breaking, the “up”-type Higgs components have electric charges of 1 and 0, while the “down”-type Higgs components have electric charges of 0 and -1.

The Lagrangian of a class of SUSY theories that respect the gauge symmetries of the SM can now be easily constructed by summing over Lagrangians in the form of Eq. 1.54, one for each supermultiplet field in Tables 1.6 and 1.7 with a superpotential built from the F -terms in Eq. 1.59. However, as we shall see in the next Section (Section 1.4.5), this basic framework still needs some minimal modifications in order to fully viable as a description of Nature.

1.4.5 Phenomenological Modifications in the MSSM

Soft SUSY breaking

Although the prediction of mass degeneracy between particles and their superpartners, demonstrated in Section 1.4.3, is of significant interest (and indeed part of why SUSY

is considered so compelling, as we shall see in Section 1.5.1), it does raise a fundamental issue when we try to write down a SUSY-enabled Lagrangian for the SM: obviously, none of the fermions in the SM have corresponding scalars with the same mass. Thus, SUSY cannot be an exact symmetry of Nature.

On the other hand, something similar also happens in the SM – the Lagrangian is symmetric under the full gauge group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$, yet the vacuum state is not, which is how interactions that have the same effect as a mass term get added, and the effective Lagrangian breaks the symmetry. In the same vein, what we know is only that SUSY is a broken symmetry of Nature at the electroweak scale; it is entirely plausible that SUSY is an exact symmetry of the full Lagrangian, but there is some mechanism that breaks the symmetry of the vacuum and thus the symmetry of the effective Lagrangian.

In practical terms, we simply consider the phenomenology of the effective Lagrangian close to the scale of the SM masses and interpret all results in terms of the parameters of such a Lagrangian. Such a treatment allows us to be agnostic towards the exact mechanism of SUSY breaking, and leaves open the possibility that at large enough scales, SUSY is an exact symmetry (which, as we shall see in Section 1.5, is needed to allow SUSY to solve several issues with the SM). This can be easily accommodated in a wide range of possible SUSY breaking mechanisms. In that spirit, we consider the most general SUSY-breaking terms in the effective Lagrangian:

$$\begin{aligned}
\mathcal{L}_{\text{soft}}^{\text{SUSY}} = & -\frac{1}{2} \left(M_{\tilde{B}} \tilde{B} \tilde{B} + M_{\tilde{W}} \tilde{W} \tilde{W} + M_{\tilde{g}} \tilde{g} \tilde{g} \right) \\
& - \left(\tilde{u} A_u \tilde{q} h_u - \tilde{d} A_d \tilde{q} h_d - \tilde{e} A_e \tilde{\ell} h_d + \text{c.c.} \right) \\
& - \tilde{q}^\dagger M_{\tilde{q}}^2 \tilde{q} - \tilde{\ell}^\dagger M_{\tilde{\ell}}^2 \tilde{\ell} - \tilde{u}^\dagger M_{\tilde{u}}^2 \tilde{u}^\dagger - \tilde{d}^\dagger M_{\tilde{d}}^2 \tilde{d}^\dagger - \tilde{e}^\dagger M_{\tilde{e}}^2 \tilde{e}^\dagger \\
& - M_{h_u}^2 h_u^* h_u - M_{h_d}^2 h_d^* h_d - (b h_u h_d + \text{c.c.})
\end{aligned} \tag{1.60}$$

where we have expressed the superfields in terms of their spinor, scalar, and vector components for clarity, and suppressed the gauge indices to avoid clutter. Here, $M_{\tilde{B}}$, $M_{\tilde{W}}$, and $M_{\tilde{g}}$ are the masses of the gauginos; A_u , A_d , and A_e are 3×3 matrices in generation-space that contain scalar (Yukawa-type) couplings for squarks and sleptons; $M_{\tilde{q}}^2$, $M_{\tilde{\ell}}^2$, $M_{\tilde{u}}^2$, $M_{\tilde{d}}^2$, and $M_{\tilde{e}}^2$ are direct mass terms for squarks and sleptons; $M_{h_u}^2$ and $M_{h_d}^2$ are direct mass terms for the up-type and down-type Higgs fields; and b is a generic cross-interaction term between the two Higgs fields. The superpotential in Eq. 1.59 gives a mass μ to the Higgsinos, while quarks and leptons of the SM still pick up their masses from electroweak symmetry breaking due to the VEV of both Higgs fields.

Naturally, this introduces a very large number of new parameters into the theory, and variations in the absolute and relative values would be expected to create a variety

of phenomenological signatures. However, in the simplest models, we assume that SUSY is broken by a single common messenger interaction that sets the scale for SUSY breaking (M_{SUSY}):

$$\begin{aligned} M_{\tilde{B}}, M_{\tilde{W}}, M_{\tilde{g}}, A_u, A_d, A_e &\sim M_{\text{SUSY}} \\ M_{\tilde{q}}^2, M_{\tilde{\ell}}^2, M_{\tilde{u}}^2, M_{\tilde{d}}^2, M_{\tilde{e}}^2, M_{\tilde{h}_u}^2, M_{\tilde{h}_d}^2, b &\sim M_{\text{SUSY}}^2 \end{aligned} \quad (1.61)$$

If we are to retain the advantages of naturalness (the strongest argument in favor of SUSY), it is clear that M_{SUSY} cannot be too much larger than the Higgs mass. This is why colliders such as the LHC (which produces collisions with $\sqrt{s}$ around that range) are extremely interesting for SUSY phenomenology.

The approach taken by experiments at the Large Hadron Collider is to interpret data in the context of various simplified signal models where the cross sections for production are standardized across analyses (as far as is possible), in order to make phenomenological interpretation as easy as feasible. The descriptions of simplified models used in this Thesis (di-gluino and di-squark production) and the assumptions that go into calculating the cross sections can be found in [7].

***R*-parity**

The SM has an appealing property: either all allowed terms in the full Lagrangian with the right mass dimension are nonzero, or there is a justification for such terms to be zero (for example, a symmetry principle). This is, however, not the case with the MSSM.

The terms appearing in the superpotential in Eq. 1.59 are in fact not the only possible gauge-invariant holomorphic products of chiral superfields. The most general superpotential would contain the following terms in addition:

$$\begin{aligned} \mathcal{W}_1 &= C_{IJK}^1 L_I L_J E_K + C_{IJK}^2 L_I Q_J D_K + C_I^3 L_I H_u \\ \mathcal{W}_2 &= C_{IJK}^4 U_I D_J D_K \end{aligned} \quad (1.62)$$

where all the C^i terms are constants. Nonzero values for any of the C^i 's would be phenomenologically quite problematic: any of the interactions in $\mathcal{W}_1$ would cause lepton number violation, while the interaction in $\mathcal{W}_2$ would cause baryon number violation. Among many other issues, these interactions would make the proton highly unstable: with natural values for the C couplings, its decay time would be much smaller than 1 s, which is obviously inconsistent with observation.

In the MSSM, this is resolved, albeit not entirely satisfactorily, by postulating that all interactions obey a new phenomenologically inspired conservation law known

as “ R -parity conservation”. Here, the following quantum number is assigned to all particles⁵:

$$P_R := (-1)^{3(B-L)+2s} \quad (1.63)$$

where s is the spin of the particle. This definition has the advantage of being very easy to evaluate in practice for all particles: it is $+1$ for all SM particles, and -1 for all their superpartners in the MSSM (that is, -1 for gauginos, sleptons, squarks, and Higgsinos). Particles with $P_R = +1$ are said to be “ R -parity even”, while those with $P_R = -1$ are said to be “ R -parity odd”. R -parity odd particles are also called “sparticles”. The statement that R -parity is conserved is the statement that a hypothesized interaction vertex is only allowed if the product of P_R for all particles in the vertex is $+1$. In other words, gauginos, sleptons, squarks, and Higgsinos can only be produced or annihilated in even numbers. This would rule out all interactions in Eq. 1.62.

This conservation law turns out to be crucial for the viability of SUSY models, and is especially important in collider phenomenology. It predicts that SUSY particles can only ever be produced in pairs (so, in proton-proton collisions, one could have di-gluino or di-squark production, but not the production of a single sparticle). Furthermore, it predicts that the lightest particle with $P_R = -1$, called the “lightest supersymmetric particle” (or LSP), is stable, because it cannot decay into other SUSY particles (being the lightest one) and cannot decay into SM particles (because that would violate R -parity). Such particles are excellent candidates for cold dark matter.

1.5 Motivation for SUSY

This Section discusses the features of SUSY that make it very attractive as an extension to the SM.

1.5.1 Solution to the Hierarchy Problem

If SUSY is indeed realized in Nature, several new fields and their interactions have to be added to the Higgs field. In particular, for every fermionic interaction of the form $\lambda_f H \bar{f} f$, a superpartner $\tilde{S}$, which is a complex spin-0 scalar, will contribute another interaction of the form $\lambda_{\tilde{S}} |H|^2 |\tilde{S}|^2$. This means that the Feynman diagram in Fig. 1.3 also contributes to the Higgs mass correction term:

Putting together the diagrams in Figs. 1.1 and 1.3, and counting contributions from both scalar degrees of freedom, the Higgs mass correction can be written as:

⁵An alternative is to define “matter parity”, a property of superfields rather than individual superfield components: $P_M := (-1)^{3(B-L)}$. Requiring $P_M = 1$ for all terms in the superpotential is phenomenologically exactly equivalent to R -parity conservation.




Figure 1.3: Correction term for the Higgs mass for a heavy complex scalar field.

$$\Delta m_H^2 = -\frac{|\lambda_f|^2}{8\pi^2}\Lambda_{\text{UV}}^2 + 2 \times \frac{\lambda_{\tilde{S}}}{16\pi^2}\Lambda_{\text{UV}}^2 + \dots \quad (1.64)$$

We now note that because particles and their superpartners have exactly the same mass, we have $\lambda_{\tilde{S}} = |\lambda_f|^2$. Therefore, it can be seen that the terms quadratic in Λ_{UV} cancel exactly, thus solving the hierarchy problem! If SUSY were realized in Nature, then the cancellation of one part in 10^{30} is not a coincidence; it would simply be an unremarkable observation that follows simply from the structure of a Lagrangian that is invariant under SUSY.

Now, we do know that SUSY is not an exact symmetry of Nature. In order for it to remain a viable solution to the hierarchy problem, SUSY breaking has to take place close enough to the electroweak scale that SUSY is recovered as an exact symmetry of Nature at high energies. Thus, the viability of SUSY as a solution to the hierarchy problem falls off quite sharply near the TeV scale.

1.5.2 Gauge Coupling Unification

By introducing a few more fields, the running of the gauge couplings in the SM is now modified. Fig. 1.4 shows the leading order calculations for the running of the gauge coupling in the SM as well as in the MSSM for various different particle masses (corresponding to the multiple solid lines). As shown, in the MSSM, the gauge couplings do unify at high energies, thus allowing for a potential Grand Unified Theory. Interestingly, for sparticle masses much above the TeV scale, the gauge couplings do not unify quite as well, which means that the LHC can probe a significant portion of the interesting phase-space for SUSY.

1.5.3 The WIMP Miracle

Whatever the particle nature of dark matter is, these particles would have been produced thermally, in abundance, in the early Universe just after the Big Bang. After




Figure 1.4: Running of the gauge couplings to leading order. SM (dashed lines), MSSM with various parameters (solid lines). Credit: [8]

some time, as the universe cooled, these particles would have decoupled from the thermal background and turned into cold dark matter.

One can show [9] that a particle with a mass around the electroweak scale (which happens to be well-probed by the LHC), if it existed, would leave behind a relic abundance that would be able to account for the entirety of the currently observed density of dark matter. Such a particle is called a “weakly interacting massive particle” (or “WIMP”). Curiously, this happens to be precisely the scale required for SUSY to correctly modify the running of the gauge couplings to allow unification at large scales (as shown in Section 1.5.2), and to solve the hierarchy problem (as shown in Section 1.5.1). Such a SUSY particle would also have exactly the sort of interactions that one might expect from a dark matter candidate (weak to nonexistent interactions with all fields of the SM). This remarkable coincidence, together with the fortunate realization that the electroweak scale is the scale of current collider experiments, is sometimes called the “WIMP Miracle”, and is a strong motivation for SUSY.

It is to be noted that there are other models for the particle nature of dark matter besides electroweak-scale WIMPs. From cosmological observations, the quantity that is well-constrained by observation is the value of $\langle\sigma v\rangle$ needed for thermal production to explain the present abundance of dark matter, and dark matter particles below the electroweak scale would also be consistent with this constraint. The requirement that the dark matter be sufficiently “cold” to allow galaxy formation at the redshift of matter-radiation equality leads to an effective floor of ~ 10 keV [10] for the mass of the dark matter particle candidate. There is a wide range of possible masses between

this floor and the electroweak scale; particles with masses below ~ 1 GeV are classified as “light dark matter”. In one such well-motivated model, dark matter candidates are pseudoscalar axions; a particular advantage of such models is their ability to explain yet another apparent SM fine-tuning, the experimentally observed absence of CP violation in QCD [11]. In fact, for a range of pseudoscalar masses, such models can also be probed in collider experiments at the LHC [12, 13].

Chapter 2

Phenomenology and Stealth SUSY

This Chapter builds upon the introduction to SUSY in Chapter 1. While Chapter 1 focuses on the theory, we describe in Section 2.1 the current status of physics searches for SUSY, and we see that large fraction of the interesting parameter-space is inconsistent with the data for many common SUSY models. In Section 2.2, we describe Stealth SUSY, one particular variant of SUSY which would be missed by many physics searches, and the motivation behind a search for Stealth SUSY in particular. The rest of this Thesis documents the methodology and the results of our search.

2.1 Current Status of SUSY searches

Given that multiple independent clues all point towards SUSY as a viable theory of Nature, it is no wonder that SUSY models have received a lot of attention at collider experiments. Typically, such models are studied starting with various simplified scenarios, some examples of which are shown in Fig. 2.1. Typically, either the neutralino or gravitino is the lightest SUSY particle and serves as the dark matter candidate. Note that in all these models, the LSP does not interact with the material of the detector and therefore creates a significant imbalance in the deposited energy. Indeed, an observed imbalance in deposited energy is a very common signature used in many analyses to distinguish SM collisions from potential collisions involving the production of SUSY particles.

Fig. 2.2 shows a compilation of recent results with limits on gluino and squark production cross sections, assuming decays to light quarks. The exact meaning of the solid and dashed lines, and the interpretation of the statistical results, is somewhat subtle and is discussed in further detail in Chapter 8, but it is reasonable to read these plots as follows: assuming some simplified models of SUSY, experiments are broadly inconsistent with regions of the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ parameter-space with $m_{\tilde{g}}$ below about 2000 GeV or regions of the $(m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ parameter-space with $m_{\tilde{q}}$ below about




Figure 2.1: Some examples of SUSY models probed by collider experiments.




Figure 2.2: Compilation of recent results with limits on gluino and squark production cross sections, assuming decays to light quarks. Credit:

<https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS>

1500 GeV.

2.2 Stealth SUSY

In Section 1.5, we saw why SUSY is considered to be a very well-motivated extension of the SM. Unfortunately, Section 2.1 seems to suggest that large portions of the meaningful phase-space, in which SUSY retains its attractiveness as a solution to various problems with the SM, are ruled out experimentally. In this Section we describe one possible resolution to the apparent tension between experimental data and the theoretical model.

2.2.1 Extensions to the MSSM with low energy imbalance

Given the numerous independent indications all hinting at SUSY as a viable theory of Nature, careful consideration is warranted before one abandons SUSY altogether. The results of Section 2.1 should in fact be interpreted as limits on no more than certain specific simplified SUSY models – in many cases, models in which there is a significant amount of imbalance in the deposited energy in a given pp collision. Given the low SM background at high energy imbalances, it does make sense to focus on such simplified models in a first study. Indeed, a significant energy imbalance is almost a trivial result of most SUSY models – considering that R -parity is conserved, the LSP is bound to be stable (in fact, it would not be a WIMP candidate otherwise), and any LSPs produced in a decay chain will not interact with the material of the detector, thus resulting in an imbalance.

But if there were a manifestation of SUSY in which there is no such significant imbalance in a collision, perhaps due to the specific kinematics of events produced with that version of SUSY, then such a model would not be ruled out by the analyses referred to in Section 2.1 – in other words, SUSY could simply be hiding in plain sight. As it turns out, there are in fact multiple SUSY scenarios that would not create any significant energy imbalance in a hypothetical collision at the LHC. These include:

- Models with compressed spectra [14, 15, 16]: in such models, there is typically a small mass gap between the LSP and the next-to-lightest supersymmetric particle (or “NLSP”). The LSP momentum is not typically small – rather, the kinematics ensures that two LSPs produced in the event have momenta in nearly opposite directions. However, such signatures are highly vulnerable to initial state radiation (ISR) giving a net boost to the collision along the beam axis. Indeed, the last study cited here uses a high missing energy trigger, so this is not truly a signature with low energy imbalance and would have been caught (albeit with lower reach) by standard analyses utilizing a high energy imbalance signature.

- Models with R -parity violation [17, 18, 19]: observed limits on the lifetime of the proton are certainly inconsistent with simple SUSY models violating R -parity. It is possible to reconcile R -parity violation in more complex SUSY models with the lifetime of the proton, but this is not very well-motivated.
- Models with hidden valleys [20, 21, 22]: in such models, TeV-scale SUSY particles play the role of mediating an interaction between the SM and a “hidden valley” (called as such because it involves a hidden sector on the other side of the “mountain”, representing the SUSY particle). In such cases the new hidden sector particles decay back into the SM, leaving no overall energy imbalance. But such models require some degree of fine-tuning in order to produce clear signatures that are detectable in collider experiments.
- Various other models [23, 24, 25, 26, 27]: all such models have some features that sacrifice some model simplicity in exchange for suppressing the energy imbalance.

In this Thesis, we shall focus on one class of models that fall under the label of “Stealth Supersymmetry” (or “Stealth SUSY”) [28, 29, 30].

2.2.2 Brief Description of Stealth SUSY




Figure 2.3: Diagram representing the key ingredients in models with Stealth SUSY. Credit: [29].

Ingredients common to all Stealth SUSY models are illustrated in Fig. 2.3.

1. In addition to the MSSM, there is a new “hidden sector” of particles at the weak scale. In the simplest Stealth scenarios, this could simply be a singlet state and its superpartner (a singlino) at the weak scale, which is the scenario considered in this Thesis.

2. As shown in the illustration, there needs to be a portal connecting the hidden sector to the MSSM – an ε -coupling which allows the hidden sector to decay to SM particles.

The essential idea in Stealth SUSY models is the following. We know that SUSY is broken in the MSSM sector. Without the presence of a portal, it would in fact be an exact symmetry in the hidden sector, but of course the sector would be completely inaccessible in that case. With the portal, SUSY breaking is communicated to the hidden sector through the ε -coupling. Because this is a weak coupling, SUSY is still approximately unbroken in the hidden sector, and therefore hidden sector particles and their superpartners are nearly mass-degenerate. The analysis presented in this Thesis is not particularly sensitive to the dynamics of the portal – all we need to assume is that one exists.

2.2.3 Motivation for Stealth SUSY

It is crucial to note that the small set of assumptions in the list above – the presence of a hidden sector at the weak scale, and the presence of a portal – is really all that is needed to hide SUSY from usual searches with a large imbalance of energy. This will be made clearer in the following paragraphs that describe the specific models in the context of which we interpret the data. We suggest that this parsimony (in terms of additional assumptions), in comparison to other models of SUSY, makes Stealth SUSY particularly interesting to study as a potential extension of the MSSM.

Fig. 2.4 shows the two simplified models of Stealth SUSY considered in this Thesis. Here, the R -parity-even state is assumed to be a 90 GeV spin-0 singlet S , while its R -parity-odd superpartner is assumed to be a 100 GeV singlino $\tilde{S}$. Again, it is important to note that the small mass splitting (10 GeV in this case) is not due to a fine-tuning that needs to be imposed on the theory – instead, it is a direct consequence of the fact that SUSY is only weakly broken in the hidden sector.

In both models, a pp collision results in strongly produced SUSY particles, with either pairs of squarks or pairs of gluinos produced in the primary interaction, as shown in Figure 2.4. These colored SUSY particles decay via the emission of quarks to neutralinos $\tilde{\chi}_1^0$, which decay to the stealth sector singlino $\tilde{S}$ via the emission of photons. The singlino then decays to its hidden sector superpartner, the singlet S , emitting the LSP, taken to be a light gravitino $\tilde{G}$. Finally, each singlet S decays to two SM gluons. The small mass splitting between S and $\tilde{S}$ naturally greatly limits the kinematic phase space available to the gravitino, which is therefore produced with very low energy. Thus the final state in both models consists of two photons, many jets, and small imbalance due to the light, soft gravitinos.

Such models have numerous advantages over other extensions to the MSSM. Note, in particular, the following features of Stealth SUSY:




Figure 2.4: Diagrams of the simplified models considered in this Thesis. The decay chain proceeds from the production of either gluino pairs (top) or squark pairs (bottom).

1. In general, there is no special fine-tuning required in these models. The small mass splitting that is essential for low energy imbalance is attributable to SUSY itself.
2. Events genuinely have low imbalance in the final state – as such, these signatures are robust against ISR boosts along the beam axis.
3. R -parity is entirely conserved in these models.
4. The LSP (the gravitino $\tilde{G}$ in both models) is a good WIMP candidate.

2.2.4 Previous Searches for Stealth SUSY




Figure 2.5: Exclusion limits on $\tilde{q}\tilde{q}$ with data from Run 1 of the LHC. Credit: [31].

Stealth SUSY has been studied before at the LHC. In particular, the parameter-space with di-squark production was probed with data from Run 1 of the LHC [31]. Exclusion limits from this analysis are shown in Fig. 2.5. However, the analysis in [31] relies

on an assumption known as S_T scaling, which we shall describe further in Chapter 6. Some recent and more reliable simulations of the major contributions to the SM background indicate the necessity of relaxing that particular assumption. This Thesis describes our significant enhancements to the background modeling procedure. We now relax the assumption of S_T scaling and instead use a simulation-based procedure to correct the data-driven background. However, despite making fewer assumptions, we still keep the systematic uncertainties under control, and we are able to exploit the increased integrated luminosity to its fullest and achieve a significant increase in the reach of the exclusion limits.

Chapter 3

Detector Description

This Chapter provides a brief description of the experimental setup. In Section 3.1 we present an overview of the full accelerator complex, and its most important component, the LHC ring. Next, in Section 3.2, we provide some details on the design and operation of the CMS experiment. Finally, in Section 3.3, we describe how the raw detector data is reconstructed and mapped to physics objects that can be used in an analysis.

3.1 The LHC

The Large Hadron Collider (LHC) is the most powerful particle collider ever constructed. It is designed to accelerate protons in two counterrotating beams to an energy of 7 TeV, resulting in a center-of-mass collision energy of $\sqrt{s} = 14$ TeV, though so far all data has been collected at or below $\sqrt{s} = 13$ TeV. It can also accelerate heavy ions to an energy per nucleon of about 2.5 TeV. The LHC is, in fact, only the last step in a chain, all of whose components have to function correctly in order to obtain collisions within the LHC. The full accelerator complex is depicted in Fig. 3.1, and a schematic layout of the LHC is in Fig. 3.2.

3.1.1 Overview of the accelerator complex

The LHC is housed in a nearly circular tunnel with circumference 26.7 km, which was originally built for the Large Electron Positron (LEP) collider. The LEP collider operated successfully from 1989 to 2000 before being dismantled to accommodate the LHC. The accelerator program operates under the umbrella of CERN, the European Laboratory for Particle Physics (the acronym derives from an older name). The main CERN site is situated on the border between France and Switzerland, and major facilities and buildings are distributed across both sides of the border. The LHC is the largest part of the accelerator complex, and its tunnel is below the surface, at a slight

The CERN accelerator complex *Complexe des accélérateurs du CERN*




Figure 3.1: Schematic layout of the full accelerator complex.

Retrieved from: <https://cds.cern.ch/record/2800984/files/CCC-v2022.png>

downward slope in the direction of Lake Geneva to accommodate the local topography. The subsurface depth of the tunnel varies from 50 m to 175 m around the ring.

As shown in Fig. 3.2, underground caverns are dug out at various points around the ring to house the experiments and accelerator facilities. There are major experiments at all four interaction points (points at which the two counterrotating beams are made to intersect – these are the points at which collisions occur). There are two general-purpose experiments: the interaction point furthest from the main CERN site is used by the CMS experiment (“Compact Muon Solenoid”), while the interaction point closest to the main CERN site is used by the ATLAS experiment (“A Toroidal LHC ApparatuS”). The other two interaction points are used by more specialized experiments: ALICE (“A Large Ion Collider Experiment”) is designed primarily to study the quark-gluon plasma, while LHCb is optimized to make precision measurements of events in which




Figure 3.2: Schematic layout of the LHC. Credit: [32]

b quarks are produced. While operation with heavy ions is a crucial component of the overall LHC physics program, we shall focus on proton-proton collisions in the rest of this Chapter and in this Thesis, because data from heavy ion collisions is not considered in the search for Stealth SUSY.

3.1.2 Accelerator Physics

The LHC is a synchrotron. Protons are accelerated with RF cavities and kept on a circular path by means of a magnetic field that, at each moment, is tuned to the energy of the circulating protons. There are two key equations governing the operation of any circular accelerator. The first one relates the required magnetic field B to the bending radius R ¹, and the momentum p and charge q of each particle. Using

¹Because the LHC is not perfectly circular, R can vary significantly across the whole ring; the equations in this Section should be regarded as first approximations meant to illustrate the general

conventional units temporarily, we have:

$$B = \frac{p}{qcR} \propto \frac{E_{\text{beam}}}{qR} \quad (3.1)$$

where c is the speed of light and the beam energy E_{beam} is the energy of each particle. All accelerating particles lose energy through synchrotron radiation, and the second equation describes the rate of this energy loss:

$$P_{\text{synch. loss}} \propto \frac{q^2 E_{\text{beam}}^4}{m^4 R} \quad (3.2)$$

where m is the mass of each particle.

3.1.3 Beam Control

Various instruments are set up all around the LHC ring to achieve specific beam characteristics. Naturally, given the size and scope of the LHC, there are many aspects of its operation that cannot be covered in a short document; nonetheless, here are the primary means by which control over the beams is maintained:

- **Resonant RF Cavities:** these use oscillating electric fields maintaining a potential difference of up to 2 MV each, to accelerate protons to high energies. The LHC has 8 RF cavities per beam that oscillate at 400 MHz. Protons in the accelerator travel in clusters, or “bunches”, that are fairly well-separated from other bunches, and the phase of the RF oscillations is finely tuned so that, as a bunch passes through the cavity, the slowest particles receive the greatest kick from the electric field. This helps maintain the bunch structure as long as there are any protons in the beam.
- **Dipole magnets:** There are 1232 superconducting dipole magnets located all around the LHC ring. The physical limitations of these magnets, in particular the maximum achievable dipole magnetic field, are the primary constraint on the beam energy of the LHC; as a result, a significant amount of attention has gone into optimizing the design and ensuring the efficient and reliable operation of these magnets. Current is run through superconducting niobium-titanium strands arranged in a solenoid pattern to achieve a magnetic field of up to 8.3 T perpendicular to the beampipes, and the whole dipole assembly is maintained, using liquid helium, at a temperature of only 1.9 K in order to maintain superconductivity and allow the strands to carry currents of up to 11 kA.

principle.

- **Quadrupole magnets:** There are a few quadrupole magnets placed around the ring, just upstream and downstream of each interaction point. These are used to “squeeze” the beam in the transverse direction by an order of magnitude just before collision. This greatly increases the probability of interaction between particles of both beams, and thus increases the luminosity at the interaction point. Quadrupoles elsewhere around the ring help to keep it focused in the transverse direction.
- **Collimators:** These are also located close to the interaction points, and their purpose is to get rid of stray particles that are too far away from the nominal beam axis, so that particles with unknown characteristics do not pollute the data in any collision event.

3.1.4 Injection Sequence

The protons that are circulated through the LHC start out in the form of hydrogen gas stored in a small cylinder – a few kg is enough to supply the LHC for months. The following sequence describes their journey from the hydrogen cylinder to the circulating beams in the LHC. Each step below is also depicted in Fig. 3.1.

- **Linac2:** Hydrogen gas is first ionized forming negative hydrogen ions, and fed to a linear accelerator that uses RF cavities to accelerate these ions to an energy of 50 MeV. (Linac2 was decommissioned in favor of an upgraded linear accelerator called Linac4 in 2018, which is capable of accelerating the ions to 160 MeV.)
- **PSB (Proton Synchrotron Booster):** The hydrogen ions from Linac2/Linac4 are stripped of their electrons and four superimposed synchrotron rings accelerate them to an energy of 1.4 GeV.
- **PS (Proton Synchrotron):** Protons from the PSB are fed into a larger synchrotron called PS, which for a few years was the world’s largest synchrotron and the earliest large synchrotron at CERN. Since the 1970s, the PS has been used primarily to feed larger accelerators, which is also its purpose in the LHC acceleration chain: it accelerates these protons to an energy of 26 GeV.
- **SPS (Super Proton Synchrotron):** Protons from the PS are fed into SPS, which is also a large synchrotron. SPS brings these protons to an energy of 450 GeV. The SPS has an illustrious past – it was used to discover the W and Z bosons – and it is still used for active research. In addition to feeding the LHC, it also provides the beams for multiple different experiments located at CERN’s Prévessin site, some of which are used to design, characterize, and calibrate the performance of future detectors.
- **LHC:** Finally, protons from SPS are fed into the LHC which ramps up the energy of each proton from 450 GeV to 6.5 TeV.

As described earlier, protons circulate in bunches. A large number of protons is injected into the LHC as described above and then the energy is slowly ramped up to 6.5 TeV. Finally, the beams are squeezed by quadrupole magnets at their interaction points and collisions can begin. Over a period of a few hours, the beams start decaying in intensity; finally, when the luminosity is low enough, the beams are all directed out of the accelerator ring and all their energy is dumped into a large block, called the “beam dump”. This whole operational duration of a few hours is called an LHC fill. There are up to 2808 bunches per beam containing about 10^{11} protons each; depending on the bunch filling scheme, only a subset of these bunches might be occupied by protons in a given fill. The design parameters of the RF cavities, dipole magnets and so on are such that bunches from opposing beams cross at each interaction point at a frequency of 40 MHz. In each bunch crossing, there are typically 40 or so individual proton-proton collisions, the vast majority of which are not very interesting for most analyses and need to be treated as background. In each event, pp collisions other than the primary pp interaction of interest are referred to as “pileup”.

3.1.5 LHC Run Schedules and Physics Performance




Figure 3.3: Integrated luminosity accumulated by the LHC over its operation, from 2011 to 2022. In this Thesis we only use collisions recorded between 2016 and 2018. Credit:

<https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>

The LHC typically operates in periodic spurts of data-taking: collisions occur for a few years, followed by a period of a couple of years in which all experiments as well as the LHC accelerator experts use the time to perform repairs and upgrades on individual subsystems. Each such spurt of data-taking is called a “Run”. So far there have been two complete runs (Run 1, in 2011 and 2012, and Run 2, in 2016, 2017, and 2018), and the LHC has recently started Run 3 in 2022. The total integrated luminosity delivered as a function of time in each year is shown in Fig. 3.3. This Thesis uses data gathered in Run 2 of the LHC.

3.2 The CMS Detector

3.2.1 Overview of the Detector




Figure 3.4: A cutaway illustration of the CMS Detector.

Credit: <https://cds.cern.ch/record/2665537>

The Compact Muon Solenoid (CMS) detector [33] is a hermetic cylindrical general-purpose detector that, with its sister experiment ATLAS, is built to probe the SM primarily near the electroweak scale as well as to enable searches for a wide variety of

potential extensions to the SM. Although both CMS and ATLAS typically probe the same physics signatures, they are built independently to different designs, with different constraints and tradeoffs, and consequently have somewhat different strengths.

A schematic of the entire CMS detector is shown in Fig. 3.4. The CMS detector, as the name suggests, is optimized for good muon identification. This requires a strong magnetic field, which is generated by the eponymous solenoid, in order to achieve good tracking performance. However, the design strength of the magnetic field constrains the size of the magnet (which explains the “compact” in CMS); and the large material budget of the steel return yoke means that all calorimetry has to happen within the compact volume of the magnet in order to achieve reasonable energy resolution. This requirement on compactness and compatibility with a very strong magnetic field in turn constrains the design of the photodetectors and other electronics used in calorimetry. Beginning with an overview of the CMS coordinate system, the next few Sections offer some more details on the design of each individual subsystem of the detector.

3.2.2 The CMS coordinate system

The coordinate system used for all work related to the CMS detector is standardized, and this is the coordinate system used throughout the rest of this Thesis. Cylindrical coordinates are used because they are well-suited for the geometry of CMS. The origin of the coordinate system is the interaction point, which is at the geometric center of the detector. The x -axis direction is defined such that it points towards the geometric center of the LHC ring. The y -axis direction is defined as the upward direction. This means that the z -axis points along the beam in the counterclockwise direction viewing the LHC ring from above. The cylindrical coordinates r , ϕ and z are defined as usual ($r := \sqrt{x^2 + y^2}$, $\phi := \arctan \frac{y}{x}$). It is also useful to define the polar angle $\theta := \arctan \frac{r}{z}$. The xy -plane is the transverse plane.

While momentum conservation obviously holds along all three axes, the colliding protons are composite particles. Along the longitudinal axis ($\hat{z}$), those particular partons whose collision is of primary interest in any given pp collision carry a poorly constrained fraction of the longitudinal momentum of their corresponding protons. Therefore, in the interests of interpretability, LHC analyses use the following convenient conventions:

- Momenta of all physics objects are projected on to the transverse plane, and the notation used for the transverse component is $\vec{p}_T$ (or just p_T for the magnitude). As mentioned above, while longitudinal momentum conservation is not very useful in hadron colliders, $\vec{p}_T$ conservation is expected to be a lot less sensitive to Lorentz boosts along the longitudinal direction.
- One particular combination of momenta is particularly useful for many analyses:

the sum of $\vec{p}_T$ over all reconstructed particles. The vector negative of this sum is referred to as the missing p_T , or $\vec{p}_T^{\text{miss}}$, and its magnitude is denoted by p_T^{miss} .

- The transverse energy E_T is defined by analogy with p_T : $E_T := E \sin \theta$.
- Often, instead of directly using the polar angle θ (or the cylindrical coordinates r and z), it is useful to express results or geometric diagrams in terms of the pseudorapidity η , defined by the equation $\eta := -\log \tan \frac{\theta}{2}$. The advantage of using the pseudorapidity η is that differences in η are invariant with respect to an arbitrary Lorentz boost along the longitudinal direction. If, for example, we have an particle produced in an LHC collision and it decays, its decay products will have the same difference in η (but not the same difference in θ) regardless of the longitudinal momenta of the incoming partons.

3.2.3 The CMS magnet




Figure 3.5: Strength (left) and field lines (right) of the magnetic field generated by the solenoid in running conditions, as predicted by a computational model. Credit: [34].

The CMS magnet is the central design feature of the CMS detector, around which the whole detector was designed. The requirement on good p_T resolution of muons and other charged particles necessitates a very strong magnetic field. In the CMS design, this field is generated by a superconducting solenoid. Many of the design decisions in the Sections below are borne out of the constraints of the solenoid design.

The solenoid is 6 m in diameter, 12.5 m in length, and 200 t in mass. Current is carried

by a 4-layer winding made from a stabilized reinforced NbTi conductor in a pure aluminum insert. The whole structure is reinforced with a light aluminum alloy and is designed to be self-supporting at full field. During regular operation, the solenoid is cooled by liquid helium to a temperature of 4.5 K (well below the superconductor critical temperature of NbTi, about 7 K at full field). The current drawn during regular operation is about 18.1 kA, which allows it to generate a magnetic field of 3.8 T in the interior of the solenoid. The predicted value of the magnetic field throughout the volume of the detector, as generated by a computational model, is shown in Fig. 3.5. This computational model has been validated through direct measurements and through regular monitoring of cosmic ray muons.

Surrounding the magnet is a massive steel return yoke, divided into multiple sections, also shown in Fig. 3.5. The barrel section of the yoke consists of five wheels, while the endcap has two disks on each end. The total mass of iron contained in the return yoke is about 10 000 t, which is a large fraction of the mass of the whole detector.

3.2.4 Inner Trackers




Figure 3.6: A cutaway illustration of the upper half of the inner tracking system. The detector is symmetric about the horizontal plane. Credit: [35].

Legend → TIB: Tracker Inner Barrel, TID: Tracker Inner Disk, TOB: Tracker Outer Barrel, TEC: Tracker EndCap.

The inner tracker system is tasked with two complementary goals:

- Providing an accurate measurement of the trajectories of charged particles. This can be used to reconstruct the energies of charged particles, and because this measurement is entirely independent of the calorimetry measurement, it is invaluable in calibrating the latter.
- Reconstructing collision vertices. This is especially important given the high pileup conditions in which CMS operates. Except in analyses with displaced

vertices and so on, physics objects that do not originate from the primary vertex can be attributed to pileup, allowing a much cleaner reconstruction of an event signature.

At peak design luminosity, there are roughly 1000 charged particles that need to be individually tracked every 25 ns. This imposes stringent requirements on the position resolution of the detector and necessitates a fast response time. The power requirements for such operation mean that the whole inner tracker volume has to be cooled. Furthermore, all detector elements have to be able to withstand the high radiation environment next to the interaction point. CMS manages these requirements with a design based entirely on silicon tracking. In a silicon tracker, a bias voltage is provided across a silicon module with a semiconductor junction to create a depletion volume. Charged particles that cross the depletion volume generate electron-hole pairs, which can be detected electronically.

Fig. 3.6 shows a cutaway illustration of the upper half of the CMS tracking system. The inner tracker has pseudorapidity coverage up to about $\eta = 2.5$. There are two primary silicon-based inner tracking subdetectors, covered below.

Pixel Tracker

The Pixel detectors are the ones that are closest to the interaction region. A recent upgrade [36] somewhat modified the geometry of the Pixel detector, so that it now has four cylindrical layers at fixed r (“barrel pixel”, or BPIX) capped by three disks each on both sides at nearly fixed z (“forward pixel”, or FPIX). The four cylindrical layers of BPIX are at $r = 29$ mm, 68 mm, 109 mm, and 160 mm, each between $z = \pm 270$ mm; while the three disks are approximately at $z = \pm 320$ mm, $z = \pm 400$ mm, and $z = \pm 495$ mm, each between $r = 45$ mm and $r = 160$ mm.

All these layers are made out of segmented silicon sensor modules, of which there are 1856, with 1184 in BPIX and 672 in FPIX. Each module has an $18.6 \text{ mm} \times 66.6 \text{ mm}$ sensor bump-bonded to an array of 2×8 readout chips (ROCs). Each ROC has 4160 readout channels, each of which is responsible for one pixel. The size of each pixel is $100 \times 150 \mu\text{m}^2$. In total, there are about 124 million readout channels in the Pixel tracker.

Particles actually cross each pixel at an angle, allowing the charge accumulation to occur over multiple neighboring pixels; an interpolation algorithm is then used to reconstruct the position of the incoming particle to a precision much better than the pixel size. This design enables a position resolution of around $20 \mu\text{m}$ in both BPIX and FPIX.

Strip Tracker

If one were to target the position granularity of the Pixel tracker throughout the inner tracking volume, the required number of readout channels would overwhelm any reasonable data acquisition system. The reason we typically need such granularity is for searches that are particularly sensitive to displaced vertices, for example b -physics in the SM, and to discriminate between pileup and the primary interaction of interest. This purpose is in fact already served to some extent by the Pixel detector. On the other hand, the pixel detector by itself is too small to provide enough information on the trajectory of a charged particle, including specifically degree to which its trajectory bends in the magnetic field, which is crucial for providing a good independent estimate of its momentum. As it turns out, because the CMS magnetic field points along $\hat{z}$ and particles generally travel along $\hat{r}$, trajectories bend along $\hat{\phi}$; therefore, what we really need in order to have a good trajectory measurement is granularity along ϕ , not so much in the other directions.

The twin requirements of a sustainable event size and good resolution in ϕ are met by designing each sensor as a long strip, with the narrow dimension (“pitch”) generally perpendicular to the ϕ direction, giving the Strip Tracker its name. There are three primary subsystems:

- TIB/TID (Tracker Inner Barrel/Tracker Inner Disk): this consists of 4 barrel layers between $r = 20$ cm and $r = 55$ cm, and 3 disks each on both ends of the barrel. The TIB/TID strip sensors are $320\text{ }\mu\text{m}$ thick. The TIB sensors have a pitch of $80\text{ }\mu\text{m}$ in layers 1 and 2, and $120\text{ }\mu\text{m}$ in layers 3 and 4, with an effective resolution of about $30\text{ }\mu\text{m}$. The TID sensors have a pitch between $100\text{ }\mu\text{m}$ and $140\text{ }\mu\text{m}$. The strips are arranged parallel to the beam in TIB, giving a precise (r, ϕ) readout; and radially in TID, giving a precise (z, ϕ) readout.
- TOB (Tracker Outer Barrel): this consists of 6 barrel layers up to $r = 116$ cm. The strip sensors are $500\text{ }\mu\text{m}$ thick, with a pitch of $183\text{ }\mu\text{m}$ in the first 4 layers (for a resolution of around $50\text{ }\mu\text{m}$) and $122\text{ }\mu\text{m}$ (for a resolution of around $35\text{ }\mu\text{m}$) on layers 5 and 6.
- TEC (Tracker EndCap): The TIB/TID and TOB together cover the region $|z| < 118$ cm. The remainder of the volume of the tracker (the range $124\text{ cm} < |z| < 282\text{ cm}$ and $22.5\text{ cm} < |r| < 113.5\text{ cm}$) is covered by the TEC subsystem. The TEC on both z -sides consists of 9 disks each, with up to 7 rings of strips with thickness $320\text{ }\mu\text{m}$ in the inner 4 rings and $500\text{ }\mu\text{m}$ in rings 5-7. The pitch varies from $97\text{ }\mu\text{m}$ to $184\text{ }\mu\text{m}$.

3.2.5 Electromagnetic Calorimetry

The inner tracker is surrounded by the electromagnetic calorimeter subsystems, whose primary purpose is to get an accurate estimate of the energy of any photons or




Figure 3.7: Cutout of one quadrant of CMS, showing ECAL and ES. (Both are symmetric about the horizontal and vertical planes.) Credit: [37].

electrons produced in a collision. CMS has two primary subsystems for electromagnetic calorimetry: the primary electromagnetic calorimeter (ECAL), and the preshower detector (ES). A cutout of one quadrant of ECAL and ES, illustrating the geometry, is shown in Fig. 3.7.

Primary electromagnetic calorimeter

ECAL is made up of several scintillating crystals, in which high-energy photons and electrons deposit most of their energy. An incoming electron interacts with a scintillator nucleus and loses some of its energy through bremsstrahlung radiation, emitting a photon. An incoming photon undergoes e^+e^- pair production in the electromagnetic field of a scintillator nucleus, and the e^+e^- pair then loses energy through bremsstrahlung in interactions with other nuclei, emitting more photons. In both cases, the energy of the incoming high-energy photons and electrons is distributed between newly produced e^+e^- pairs and bremsstrahlung photons. Each of these new particles produces yet more e^+e^- pairs and bremsstrahlung photons through the same processes, and this chain continues until the energy of the photons produced is too low to generate any more e^+e^- pairs; in this regime, ionization takes over as the primary means of energy loss. The net result of these interactions is a nearly instantaneous cascade of high-energy particles, called an electromagnetic shower. Numerous photons are produced at wavelengths at which the scintillator material is transparent. These photons can be picked up by a photodetector attached to the crystal, whose response can be calibrated to provide an accurate estimate of the energy of the incoming particle that set off the cascade. There are two material properties of particular importance in calorimetry:

- The Molière radius of a given material quantifies the transverse size of an

electromagnetic shower in that material. This is the diameter of an imaginary transverse cylinder that contains 90% of the shower. In a calorimetric application, a lower Molière radius allows for better position resolution.

- The radiation length is the characteristic distance over which an incoming particle loses a fraction $(1 - \frac{1}{e})$ of its energy. While the Molière radius quantifies the transverse size of a shower, the radiation length can be used to determine how quickly an incoming particle loses energy as a function of longitudinal depth. In a calorimetric application, a lower radiation length requires a smaller scintillator depth to achieve the same degree of energy dissipation.

The design of ECAL was guided by the constraints mentioned in Section 3.2.1. The crystals had to be compact to be accommodated inside the magnet volume, and had to have good transparency characteristics to allow the readout from a photodetector to accurately represent the energy dissipated by an incoming particle. This necessitates a dense material with heavy atoms. ECAL uses lead tungstate (PbWO_4), which has the additional advantage of a very fast response time (~ 25 ns, close to the timing between adjacent bunches), enabling good timing resolution and avoiding contamination of the readout from out-of-time hits in neighboring bunch crossings. Lead tungstate has a reasonably small Molière radius of 2.19 cm, and a short radiation length of 0.89 cm. The range of wavelengths in which lead tungstate is transparent happens to be around the visible range, which is an advantage because this allows calibration to be performed by a laser system operating around that range.

ECAL consists of a cylindrical section surrounding the central beam-pipe in the pseudorapidity range $|\eta| < 1.48$, and endcaps at both ends that extend the pseudorapidity coverage to $|\eta| < 3.0$, as shown in Fig. 3.7. The scintillating crystals have the same depth, 23 cm or about 25 radiation lengths, in the barrel (EB) and both endcaps (EE+ and EE-), and they all have photodetectors glued to the back of the crystal to pick up the scintillation signal. However, the crystals have slightly different transverse dimensions and the photodetectors are quite different between the barrel and endcaps because the operating environments are very different:

- In EB, there are 61200 crystals with transverse dimensions $2.2 \times 2.2 \text{ cm}^2$. The strong magnetic field is nearly parallel to the photosensitive surface, which means that conventional photomultipliers would not work because the strong field would disrupt the paths of the electrons in the PMT. Instead, EB uses avalanche photodiodes (APDs), which work by applying a strong bias voltage across a depletion region that is only $\sim 6 \mu\text{m}$ thick, so that the scintillation signal is nearly unaffected by the magnetic field. However, APD noise is highly temperature-sensitive, and ECAL APDs have to be kept at a temperature of $18 \pm 0.05^\circ\text{C}$. Under nominal operation these APDs achieve a gain of 50.
- In EE+ and EE-, there are 14648 crystals in total across both endcaps with transverse dimensions $2.9 \times 2.9 \text{ cm}^2$. While the magnetic field is also strong

near the endcaps, it is perpendicular to the photosensitive surface, which creates less of a problem with traditional photomultiplier tube designs. On the other hand, the radiation environment is significantly harsher than for EB, which creates a problem for APD-type designs. Accordingly, EE+ and EE- use vacuum phototriodes (VPTs), which are more radiation-tolerant. They also have the advantage of not being very sensitive to temperature fluctuations unlike the APDs. However, they do have a smaller gain than APDs – 10 in nominal operation – which is partly ameliorated by having a larger photosensitive area per crystal, made possible by the larger transverse size of the endcap crystals in comparison to the barrel crystals. This comes at the cost of a slight loss of position resolution in the endcaps.

Over time, the lead tungstate crystals as well as some of the frontend electronics degrade due to radiation damage. In order to ensure that the effect of this degradation on physics analyses is characterized and controlled, ECAL uses two primary techniques. First, a laser system continuously injects pulses of blue and green light with known intensity into each crystal, cycling over all crystals in ECAL with a periodicity of about 20 min when the detector is in operation. This allows us to characterize the time evolution of the calibration for each crystal. However, this is not a very precise process, and needs to be augmented with a more precise intercalibration technique [38] that is much more responsive to differences between crystals. This is needed to get the energy resolution of ECAL close to its design specifications. Intercalibration uses three independent “standard candle” measurements:

- Using the known mass of the pion and looking at crystal energies in $\pi^0 \rightarrow \gamma\gamma$ events.
- Matching the reconstructed energy in ECAL, E , to the value of p_T as measured by the inner tracker. The idea is that the distribution of $\frac{E}{p_T^{\text{miss}}}$ should remain stable over time in any given η bin.
- Ensuring that the $Z \rightarrow ee$ peak yields the correct mass and lifetime, known from LEP results.

Preshower

In addition to the scintillating crystals, the electromagnetic calorimeter also contains 137,216 thin silicon strips with dimensions $2.0 \times 63 \times 0.3 \text{ mm}^3$ in four planes in a sampling configuration (two planes in front of each endcap with coverage $1.65 < |\eta| < 2.6$), with lead plates of thickness $3 X_0$ as the absorber. These strips operate similarly to those used in the inner tracker.

Approximately 5% of the total energy deposited by an electromagnetic particle in the endcaps is absorbed by the preshower. The primary design purpose of the ECAL preshower is to achieve excellent spatial resolution in the endcaps, which is necessary

to distinguish high-energy photons from close photon pairs (for instance those resulting from π^0 decay). This is achieved by orienting the thin strips lengthwise along x in one plane and along y in the other plane in front of each ECAL endcap. This increases position resolution in the endcaps by an order of magnitude.

Net Energy Resolution

In all, electromagnetic calorimetry allows precise reconstruction of the energies of incoming photons and electrons. The energy resolution actually depends on the energy of the incoming particle, but for photons with $p_T \sim 60$ GeV, the resolution achieved is better than 3% for $|\eta| < 1.4$.

3.2.6 Hadronic Calorimetry




Figure 3.8: Cutout of one quadrant of HCAL.

Legend → HB: Barrel, HE: Endcap, HO: Outer, HF: Forward.

(The detector is symmetric about the horizontal and vertical planes.)

The (unlabeled) inner trackers and electromagnetic calorimeter are also shown for geometric context. Credit: [33].

ECAL is surrounded by the hadronic calorimeter (HCAL), which is designed with two complementary goals: to measure the energy of hadronic jets, and to better constrain the missing transverse momentum p_T^{miss} in each collision. A cutout of one quadrant of

HCAL, illustrating its geometry, is shown in Fig. 3.8. Typically, hadrons, muons, and any other long-lived particles produced in a collision are not stopped by ECAL – to stop such particles, what is needed is a large amount of dense material with enough nuclei to collide inelastically and interact with the incoming hadrons to dissipate their energy.

We saw that the Molière radius and the radiation length are important material characteristics for electromagnetic calorimetry; for hadronic calorimetry, the relevant material characteristic is instead the nuclear interaction length, which, for a given material, is the mean free path of a hadronic particle in that material before it interacts with a nucleus. Of course, in the path followed by any hadron, there is a significant degree of nuclear interaction between the collision point and HCAL, accounting for about 1.1 interaction lengths. This, among other reasons, lowers the energy resolution of HCAL compared to ECAL.

HCAL is divided into four components, three of which are sampling calorimeters (like ES and unlike ECAL). In a sampling calorimeter, the scintillating material is different from the absorptive material: so the calorimeter is made out of alternating layers of absorber and scintillator. The absorber dissipates the energy of the particle, and the scintillator measures how much energy the particle has as a function of travel depth. Particles with different energies have different profiles for the energy dissipated as a function of depth, so that a combined readout at multiple depths is sufficient to reconstruct the energy of the original particle. All scintillator layers in HCAL are plastic scintillators, and embedded optical fibers convey scintillation signal outside the detector where it is picked up by a hybrid photodiode (HPD).

Each of the four components is described below.

HCAL Barrel

The HCAL barrel (HB) covers the pseudorapidity range $|\eta| < 1.3$. Its radius ranges from the outside surface of the ECAL barrel to the inside surface of the CMS magnetic coil: $1.77\text{ m} < R < 2.95\text{ m}$, which is only enough to cover 5.82 interaction lengths for a particle emitted perpendicular to the beam. The HB has a total of 16 absorber and 17 scintillator layers. The absorber layers are all made from brass, except the first and last which are made from stainless steel for structural support. The layers are segmented into tiles that each cover an area of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$, roughly 25 times the area covered by a single crystal in the ECAL barrel.

HCAL Endcaps

The HCAL endcaps (HE) cap off HB on both sides and cover the range $1.3 < |\eta| < 3.0$. The thickness is about 10 interaction lengths, made up of 17 absorber and 18 scintillator layers. The segmentation of these layers corresponds to $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$

for $1.3 < |\eta| < 1.6$ and $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$ for $|\eta| > 1.6$. HE also uses brass for the absorber layers. Brass is ideal as an absorber for both HB and HE because of its nonmagnetic characteristics, useful in a region with high magnetic field.

Outer Calorimeter

The Outer Calorimeter (HO) acts as a “tail catcher” for the significant number of hadrons that are not stopped by HB. The CMS solenoid has a return yoke made out of steel, which act as the absorber for HO. The geometry is complex and constrained by the magnet design and the muon systems; in general, there are either one or two HO scintillator layers that an outgoing jet might interact with. This increases the effective total thickness of the hadronic calorimeter system to about 10 interaction lengths below $|\eta| < 1.3$.

Forward Calorimeter

In many physics analyses, the total missing transverse momentum (p_T^{miss}) plays a crucial role, and it is important for it to be measured as accurately as possible. If only hadronic jets up to $\eta = 3.0$ are reconstructed (from HB, HE and HO), that leaves open the possibility that a high-energy jet was emitted in the collision at higher η but simply missed the detector acceptance. Practically, this would lead to a high uncertainty on p_T^{miss} for each collision. To reduce the effect of this issue as far as possible, the fourth component of the hadronic calorimeter, the forward calorimeter (HF), covers the range $3.0 < |\eta| < 5.2$. The key challenge in this operating environment is the extreme radiation; accordingly, the design is kept as simple and radiation-hard as possible. HF is designed as a large mass of steel absorber, with embedded quartz fibers which detect particles via Čerenkov radiation.

3.2.7 Muon Systems

Good muon performance is a primary design goal of the CMS experiment. To achieve good muon identification and energy reconstruction, readouts from the inner tracker are combined with those from three muon subdetectors. A cutout of one quadrant of the three muon systems, illustrating their geometry, is shown in Fig. 3.9.

Any electrons or photons produced in a collision cause electromagnetic showers in ECAL, while light quarks and gluons promptly hadronize and result in hadronic showers within HCAL. Heavier particles such as tau leptons or top quarks decay mostly into the lighter particles mentioned above. Neutrinos (and potentially, any new stable particles that are not in the SM) are not designed to be detected directly in a collider experiment. The only particles left are muons. Because muons are much heavier than electrons, they only lose a small fraction of their energy as they interact with the ECAL, HCAL, and return yoke material. Instead of producing a shower,

they leave a trail of ionized particles, and punch through the whole detector. It is the singular purpose of three muon subdetectors to characterize any muons produced in a collision. Only muons reach these subdetectors, and because they are on the outside of the solenoid, the constraints on size are somewhat more relaxed compared to the inner tracker and the calorimeters.

Each of the three muon subdetectors is covered below.

Drift Tubes

The Drift Tube (or DT) muon tracker is on the outside of the solenoid and covers the pseudorapidity range $|\eta| < 1.2$. In general a drift tube works on the following principle: a wire at high positive voltage acts as an anode. Surrounding the wire is a gas. When a muon passes through the gas, it ionizes its atoms creating some free electrons and positive ions. These electrons start drifting towards the wire, and as they accelerate, they in turn ionize other atoms in their path, amplifying the number of free electrons that eventually reach the wire. The time taken by the electromagnetic cascade to reach the anode can be used to reconstruct the perpendicular distance of the muon track from the wire.

For the CMS drift tubes, the gas is a mixture of 85% Ar and 15% CO₂. The basic unit of the DT detector is a cell – a rectangular aluminum tube with dimensions 42 mm $\times$ 13 mm, which acts as the cathode, contains a single gold-plated stainless steel wire with diameter 50 μ m, which is the anode; the anode is kept at a voltage of 3.6 kV. These cells are arranged in layers either parallel to the beam (along $\hat{z}$), or perpendicular to the beam, providing a good measurement of a muon’s (r, ϕ) (or (r, z)) coordinates at some point in its trajectory. The length of the wire is 2.4 m for drift tubes along $\hat{z}$.

A set of four such layers staggered by half a cell width forms a “superlayer” (SL) and acts as one independent unit in muon reconstruction. A DT “chamber” is made of either three or two parallel SLs. Chambers are combined into DT “stations”, that are interleaved with the return yoke of the magnet. The inner three stations (MB1, MB2, MB3 in Fig. 3.9) contain 60 chambers, while the outer station contains 70 chambers. The chambers in the inner stations have one (r, z) SL placed between two (r, ϕ) SLs, while the chambers in the outer stations only have two (r, ϕ) SLs.

In total, there are 172000 cells in DT, each of which has an position resolution of about 250 μ m (for individual readouts) and a position resolution of about 100 μ m when readouts are statistically combined.

Cathode Strip Chambers

The cathode strip design, used in cathode strip chambers (CSCs), has some similarities with the drift tube design, but the principle of operation is quite different. In this

design as well, wires kept at high positive voltage acts as an anodes and are suspended in a gas. The cathodes, however, are now metallic strips that run perpendicularly to the wires. Just as in the case of drift tubes, a muon ionizes the gas, and electrons, as they accelerate towards the anode, create a cascade of other electrons and positive ions. However, here, the positive ions also play an important role – they are deposited on the cathode strips and are picked up in the readout. The width of the strips is smaller than the size of the typical shower, so that positive ions are deposited across multiple strips; the relative fractions of the charges deposited in neighboring strips allows for a precise reading of the position of the muon in the direction perpendicular to the strip. Because the wires are also perpendicular to the strip and are also picked up in the readout, this allows for a precise determination of the position along both directions perpendicular to the travel path of the muon.

CSCs have a few advantages over a drift tube design at high η : they are able to operate at much higher rates, allowing rapid muon triggering; they can operate in strong but non-uniform magnetic fields (at low η , the field is fairly uniform); and the strip design naturally lends itself to the geometry of the detector because a radial fan-shaped pattern for the strips is quite naturally compatible with the principle of operation.

For the CMS CSCs, the gas used is a mixture of 40% Ar, 50% CO₂, and 10% CF₄. In η , CSCs cover the range $0.9 < |\eta| < 2.4$ (there is some overlap with the DTs). There are four CSC station layers in each endcap. The station layer closest to the interaction point (labeled “ME1” in Fig. 3.9) has three annuli, the second and third layers (“ME2” and “ME3”) have two annuli, while the fourth layer (“ME4”) only has one annulus at high η . Each annulus is divided along ϕ in 18 or 36 trapezoidal modules. Each module contains 6 chambers, each of which has a number of radial strips segmented in ϕ (and thus able to provide a precise reading of ϕ), while the anode wires, being perpendicular to these strips, provide a precise reading of η .

In all, there are about 220000 cathode strip readout channels and about 180000 anode wire readout channels in the CSCs. The position resolution achieved along ϕ is roughly 150 μm .

Resistive Plate Chambers

DTs and CSCs are designed to provide highly accurate information on the full trajectory of a muon. However, because the delay between the passage of a muon and the avalanche reaching the anode or cathode could be more than 25 ns (the time between two consecutive bunch crossings), there is some ambiguity in the bunch crossing corresponding to the production of a given muon. This would make operating at high trigger rates very challenging. (More information on the CMS Trigger system will be provided in Section 3.2.8.) The resistive plate design features excellent timing resolution (much better than the collision frequency), and, by so

removing the ambiguity on bunch crossing, allows operation at very high trigger rates.

The fundamental element of this design is a set of two parallel plates made with a highly resistive material, with the space between the plates occupied by a gas. The plates are maintained at a voltage difference just below breakdown voltage. When an incident muon enters the gap between the plates, it ionizes the gas just as in a drift tube, but, because the voltage applied is much higher, the avalanche reaches the plates much faster. (This would create a problem if we were trying to measure position precisely as in a drift tube, but for the intended use-case of RPCs, a coarse measurement of position is perfectly acceptable.) Because the plates are highly resistive, the discharge is highly localized and ends when all the ionized gas has reached the plates. Naturally, this design also makes it difficult to read out the avalanche signal through any direct interface with the plates; instead, the signal is picked up as an image current induced on a conductive strip on the outside of the plates. The time between the arrival of the avalanche and the subsequent recharging of the plates (after the avalanche has been dissipated) is a dead zone for that particular localized portion of the plates (in the sense that any muons that happen to pass through that precise region in that time duration would not be detected), but, in practice, this does not present a significant problem at the muon flux expected in regular operation.

In CMS, the plates in the RPCs are made of 2 mm thick bakelite, the gap between the plates is 2 mm, and the gas between the plates is primarily $\text{C}_2\text{H}_2\text{F}_4$. Potential differences for each pair of plates are optimized individually; by and large, they are maintained at a potential difference of around 9.5 kV.

RPCs are installed both in the CMS barrel and in the endcaps (though coverage in the endcaps only goes up to $|\eta| < 1.6$). The barrel modules are inserted into the DT stations. There are six barrel layers: two each in the first two inner stations, and one each in the two outer stations. Along ϕ , the modules have the same segmentation as DT modules. There are three endcap layers on each side, divided into three annular layers; along ϕ , these are also segmented in trapezoidal sections just like the CSCs, and have the same $\Delta\phi$ segmentation as the barrel modules.

In all, there are 80640 readout strips in the barrel and 92162 readout strips in the endcaps.

3.2.8 Trigger Systems

The LHC produces collisions at a rate of 40 MHz. The space required to store event data varies depending on the physics of a given collision, but ~ 1 MB is a reasonable order of magnitude estimate. If every single collision were to be stored (even if that were possible given the limitations of the electronics, such as buffer readouts and so on), this would mean a bandwidth requirement of about $\sim 40 \text{ TB s}^{-1}$. The bandwidth alone would be a nearly insurmountable obstacle, not to mention the storage requirements

and the efficiency of a full analysis workflow. Luckily, the vast majority of these collisions are simply soft scattering events from pp collisions and are not of particular interest to analysts working with CMS data. This is why CMS employs a two-level filtering scheme to ensure that the limited data budget is spent on storing data only from those collisions that are interesting for at least some subset of analyses. In the first filtering step, called the L1 trigger, coarse-grained data is analyzed very fast and the rate of events to be analyzed further is brought to less than about 100 kHz. In the second filtering step, called the High Level Trigger, data at the full granularity of the detector is analyzed using a computing farm and events are selected for final storage at a rate of a few hundred Hz.

L1 Trigger




Figure 3.10: The CMS L1 trigger workflow. Credit: [40]

The Level 1 (or L1) Trigger uses coarse-grained data from the electromagnetic and hadronic calorimeters and from the muon detectors to make a near-instant decision on whether or not an event is potentially interesting enough to be forwarded to the expensive high-level trigger. The L1 trigger workflow is illustrated in Fig. 3.10. Given

the design requirements, L1 modules are firmware-based and many are located on the detector itself; the requirements on radiation hardness and calculation speed lead to the choice to perform these computations on FPGAs and ASICs. At the most local level, hits in calorimeters and the muon detectors are organized into trigger primitives (TPs).

Calorimeter TPs contain basic information about various quantities related to any EM showers or hadronic activity in an event at a resolution corresponding roughly to an area of 5×5 ECAL crystals ($\Delta\eta \times \Delta\phi = 0.087 \times 0.087$), and each such deposit is associated with its bunch crossing. These TPs are sent to a dual-step calorimeter pipeline, in which Layer 1 preprocesses and collates the TPs received from the calorimeters, and Layer 2 performs object identification and energy reconstruction. This is only a coarse-grained version of the full reconstruction, which is described in more detail in Section 3.3. Clustering algorithms are used to identify EM candidates (electron or photon) by looking for ECAL energy maxima above a threshold; a spread in deposited electromagnetic energy along ϕ is attributed to bremsstrahlung radiation from electrons and can thus be used for disambiguation between electrons and photons. In this context, it is used to apply corresponding calibration factors to get a more precise reading on the energy of the reconstructed particle. Jets are reconstructed with a 9×9 sliding window of “trigger towers” (a well-defined collection of detector elements such as crystals) centered around a maximal HCAL deposit. Missing energy is reconstructed at the full granularity of the detector. Simple algorithms then encode all this information into a series of bits, corresponding to whether or not a given collision passes the L1 criterion at that bit position, and send it forward to the global trigger.

The muon L1 workflow proceeds similarly. TPs are analyzed by all three muon detectors, which reconstruct tracks within each individual chamber; these TPs are sent to the track finder (TF) layer. The TF layer then attempts to combine these tracks (using geometric information about the positions of each chamber) into unified muon track candidates. The next layer attempts to improve the quality of reconstructed tracks by algorithmically merging nearby or overlapping tracks. This information is then combined with calorimetric information, which includes the positions and energies of any photon or electron candidates or jets in an event, and fed to a global muon L1 trigger, which computes its own bits just as in the L1 calorimeter workflow and sends these to the global trigger.

The data from each bunch crossing is saved temporarily by each subdetector in a buffer. If an L1 accept signal is sent to the subdetectors for a given bunch crossing, each of them sends subdetector information (at full granularity) via a Data Acquisition (DAQ) system to the next filtering step – the high level trigger.

High Level Trigger

Collisions that get handed to the High Level Trigger (HLT) are those events that the coarse-grained L1 algorithms were unable to reject. As such, it is expected that one would need to perform a much more robust analysis on these collision events. Thus, the algorithms are by necessity much more sophisticated than those run by the L1 trigger, and need to be very flexible (to keep up with evolving detector conditions, beam conditions, etc.). On the other hand, given the relatively low rate at which events get sent to HLT (as compared to the collision rate), not much is gained by using on-detector electronics or specialized firmware.

As a result, the high level trigger is entirely a software-based system that uses off-the-shelf computing hardware. A so-called “HLT farm” with more than 10000 CPU cores is located on site. Data from every collision that passes the L1 global trigger is analyzed using the same software framework that is used in offline analysis. There are more than 400 different filters that select for a broad spectrum of interesting physics signatures; each such filter is called a “trigger path”. For example, one of the trigger paths in the analysis presented here requires two reconstructed photons in an event with $p_T > 30$ GeV and $p_T > 18$ GeV; there are also some loose quality criteria such as a threshold on the amount of non-electromagnetic energy in a cone surrounding the reconstructed photons. Each trigger path has a well-defined subset of L1 bits that all need to be set in order for the algorithm to run that particular trigger path for a given collision. Furthermore, all these filters are designed to run well-defined sequential conditions: the first condition that is not met aborts any further processing, minimizing unnecessary computation.

Each trigger corresponds to a primary dataset; groups of triggers that share a broad physics signature category are categorized into the same primary dataset. For example, events exhibiting significant electromagnetic activity are fed into the **EGamma** dataset. The thresholds for the decision tree in each trigger are set such that the trigger rate for a given trigger is not too high, so that HLT does not exceed its total bandwidth budget of a few hundred Hz.

Triggers can be run selectively on a fraction of events. Triggers which are only active for a fraction $1/N$ of recorded events are called “prescaled” triggers and are said to have a prescale of N . Prescaled triggers are typically used as input for commissioning new triggers, for characterizing the efficiency of existing triggers, etc. Typically, the instantaneous luminosity (and therefore the rate of interesting events) is highest near the beginning of an LHC fill and gradually decreases over the duration of the fill. The prescale mix of enabled triggers is therefore changed over the duration of a fill. At the beginning of a fill, prescaled triggers are generally active with a high prescale value; over time, as the instantaneous luminosity and event rate decreases, the value of the prescale on these triggers is gradually reduced in such a way that the total trigger rate is approximately stable.

There are some assumptions that are implicit in every physics analysis that relies on a specific trigger. It is possible that even if a collision with the target signature were produced during an LHC collision, it might be missed by the L1 or HLT triggers. This could be due to one of many reasons: there might be an electronics issue either preventing readout altogether in a particular region of the detector; TPs from some region in the detector might be prevented by a link error from being sent to the next L1 stage; it might be that two interesting collisions happen in consecutive bunch crossings, which means that, given buffer limitations, only the first one is actually saved by the L1 and HLT; the event might simply not be reconstructed properly; and so on. For all these reasons, every physics analysis has to characterize the efficiency of any triggers used by the analysis – that is, the probability that an event with the target kinematics would indeed survive the L1-HLT filter.

3.3 Physics Object Reconstruction

In Chapters 1 and 2, the objects we dealt with were abstract particles and fields: electrons, gluinos, and so on. In this Chapter, so far we have detailed the process of getting an electronic readout from detector elements. We now present some details on how CMS translates a set of detector deposits in a bunch crossing into a set of physics objects for an event with accurately reconstructed properties (such as the p_T of the object and the direction in which it is emitted in a collision). This allows analysts to match events produced in experimental data to abstract Feynman diagrams or other field-theoretic constructs. This translation is carried out by an algorithm called the Particle Flow (PF) algorithm.

Finding the correct association of detector readout hits to physics objects relies on the differences in the behavior of various particles in the CMS detector, illustrated in Fig. 3.11. Muons leave deposits in the inner trackers as well as the muon chambers. Muons also deposit some energy in both calorimeters, but the energy deposited is typically quite small compared to the energy deposited by photons or jets. Electrons leave deposits in the inner tracker and deposit most of their energy in ECAL, so they leave little to no deposits in other layers; just as for the muon. Charged hadrons (like π^+) leave deposits in the inner tracker as well as ECAL, but most of their energy is deposited in HCAL. All particles we have discussed so far are charged particles (which is why they leave deposits in the inner tracker), and their charge can be inferred from the direction of the bending of the particle in the constant magnetic field produced by the solenoid. Neutral hadrons (such as K_L^0 mesons and neutrons) do not leave any deposits in the inner tracker, but they do deposit some energy in ECAL and most of their energy in HCAL. Photons do not leave any deposits in the inner tracker either, and they deposit most of their energy in ECAL.




Figure 3.11: Trajectories of various particles through the detector. Credit: [41]

3.3.1 Tracks

The inputs for tracks and vertices are readouts from the inner trackers (both the Pixel and Strip trackers). Reconstruction relies on an iterative combinatorial track finder (CTF) that uses a Kalman filter [42]:

- In a first pass, “seeds” are constructed from either triplets of neighboring hits or pairs of neighboring hits with some other quality criteria.
- A track consistent with these seeds is grown in both directions, and any matching hits are added to this track and removed from the full collection. Because the parameters of the candidate track evolve as particles radiate or scatter, an adaptive Kalman filter is used.
- These steps are repeated iteratively with increasingly sophisticated versions of seed finding in order to capture as many tracker hits as possible.
- A final fit is performed on all tracks found by the algorithm to fine-tune its parameters, and the normalized χ^2 value, impact parameter, and so on are stored

for each track.

Once a set of tracks has been determined, vertex reconstruction is carried out. For each events, a set of points is determined such that each point is at a local maximum of the likelihood of reconstructed tracks passing through it. This is a nontrivial task, especially in conditions of high pileup. An efficient algorithm based on the “deterministic annealing” technique [43] is used for this clustering and scales well with pileup. The primary vertex is determined algorithmically as the vertex with the highest $\sum p_T^2$ of associated tracks. In general this primary vertex is the one that is of interest in a physics analysis.

This procedure is able to reconstruct vertices with a resolution of about $10\text{ }\mu\text{m}$ at the pileup expected in running conditions.

3.3.2 Electrons and Photons

A clustering algorithm first runs on deposits in the electromagnetic calorimeter. First, seed crystals are identified as ECAL crystals with a local maximum of deposited energy, and seed clusters are found by including a few neighboring cells around the seed crystal. The exact number varies between the barrel and endcaps. Next, any neighboring crystals that overlap with an existing cluster are merged into the cluster. This process is repeated until all energy deposits above a certain threshold are accounted for. Bremsstrahlung emission from incoming electrons results in a spread along ϕ of energy deposits; similarly, photons often convert into e^+e^- pairs, which also produce a spread along ϕ of energy deposits. Thus, we define superclusters as collections of neighboring clusters that are spread tightly in η but may be spread loosely along ϕ .

Those superclusters that are isolated from tracks not matched to the supercluster, and isolated from other superclusters, are selected for further processing. For each such supercluster, the energy deposited in HCAL just behind the supercluster is computed and, if it happens to be significantly smaller than the energy contained in the supercluster, then the supercluster is regarded as an electron or photon candidate.

Disambiguation between electrons and photons is performed by examining the tracks reconstructed in the event. Electrons interact with the tracker, while photons (largely) do not; thus, if there is a reconstructed track leading to the position of the supercluster object, then the object is labeled an electron; otherwise the object is labeled a photon.

The reconstructed energy is then calculated as a sum over all crystals within the corresponding supercluster:

$$E_{e,\gamma} = F_{e,\gamma} \times \left[G(\eta) \times \left\{ \sum_{i \in \text{SC}} S_i(t) \times C_i \times A_i \right\} + E_{\text{preshower}} \right]. \quad (3.3)$$

Here, the factors A_i are the raw pulse amplitudes in ADC (there are ten of these,

spread in time around the bunch crossing), as measured by the read-out electronics. A detailed description of the raw amplitudes can be found in [44]. $G(\eta)$ contains the ADC-to-GeV conversion factor allowing for an η -dependence to incorporate geometric effects. The factor $F_{e,\gamma}$ corrects for imperfect clustering and other geometric effects (such as material upstream of the detector), and is different between a photon and an electron candidate. For electrons and photons incident on the endcaps, $E_{\text{preshower}}$ contains an estimate of the energy absorbed in the lead planes. The factors $S_i(t)$ are obtained from the laser response correction, while the inter-calibration coefficients C_i allow us to set an energy scale at the granularity of a single crystal; both of these were touched upon briefly in Section 3.2.5.

3.3.3 Muons

Muons leave deposits in the inner tracker as well as the muon chambers. The algorithm first reconstructs standalone muons from the muon chambers. Extrapolating the trajectories of these standalone muons to the inner tracker, all tracks that match such trajectories are classified as tracker muons. A global muon is formed by performing a joint fit of the inner and outer tracks [45].

Muons reconstructed in the chambers are typically higher purity, while momentum measurements of muons reconstructed in the inner tracker are typically more precise. Analysts choose which muon identification criteria to use for a given analysis based on this tradeoff.

3.3.4 Jets

Fundamental physics makes jet reconstruction much more difficult than muon, electron, or photon reconstruction. It is more complex to implement in practical terms, and more subtle to interpret in terms of the underlying theory.

When, for example, a Feynman diagram of a collision has a final state photon, the link between that abstract photon and a photon reconstructed as in Section 3.3.2 is straightforward to follow. However, this is not the case for final state quarks or gluons. It is reasonable to analyze and run simulations for Feynman diagrams with such final state partons in the perturbative regime, but the richness and complexity of nonperturbative long-range QCD interactions ensures that all such partons very quickly hadronize if such events are actually produced in a collider experiment. From the theory perspective, what we have in the final state is a loose collection of partons generated by hadronization of the original parton, and this collection can only be defined in stochastic terms.

The distributions of any physical observable used in an analysis have to be well-predicted by the theory, which results in certain constraints given the complexity of hadronization. In particular, because calorimetric deposits are used in order to

determine the number of jets and the p_T of each jet, any clustering algorithm that assigns a calorimetric cell to a given jet has to obey the following two criteria:

- Infrared-safety: the clustering has to be insensitive to the emission of a soft parton with $E \rightarrow 0$ (that is, adding a low-energy jet to an event should not impact the number or p_T of each reconstructed jet).
- Collinear-safety: the clustering has to be insensitive to a parton splitting into two nearly collinear partons.

CMS uses the so-called “anti- k_T ” clustering algorithm [46], which is both infrared- and collinear-safe. In this clustering algorithm, all calorimetric cells with significant energy deposits start out as “protojets” (indexed by i below) and are assigned the following characteristic value:

$$d_i := \frac{1}{p_{T_i}^2}. \quad (3.4)$$

The distance between two protojets indexed by i and j is defined as:

$$d_{ij} := \min(d_i, d_j) \times \frac{\Delta R_{ij}^2}{R^2} \quad (3.5)$$

where R is a global parameter and $\Delta R_{ij} := \sqrt{(\phi_i - \phi_j)^2 + (\eta_i - \eta_j)^2}$. In the algorithm, at any step of the iteration, d_i is calculated for all protojets and d_{ij} is calculated for all pairs of protojets, and all these values are compared to find the smallest one. If d_{ij} is the smallest among these values for some i and j , then protojets i and j are merged into a single protojet with a p_T equal to the sum. Otherwise, if d_i is the smallest among these values for some i , then protojet i is promoted to a jet and removed from further consideration. This iterative procedure is repeated until all protojets have been clustered into jets.

The parameter R can be tuned to the desired value: increasing the value of R increases the size of the cone that contributes to the energy of a given jet and thus captures more of the hadronized shower; but increasing the value of R also increases the probability of contamination by pileup. Two primary values of R are used in the standard reconstruction: $R = 0.4$ and $R = 0.8$. Throughout this analysis, we use jets that have been clustered with $R = 0.4$.

3.3.5 Missing transverse momentum p_T^{miss}

The definition of the variable p_T^{miss} , and the reason it is of particular interest in physics analyses, has already been introduced in Section 3.2.2. Once the data from all subdetectors for a given collision has been organized into a set of electrons, photons,

muons, and hadrons, and the $\vec{p}_T$ of each particle has been found, the missing transverse momentum is computed quite simply as $\vec{p}_T^{\text{miss}} := -\sum_{\text{PF candidates}} \vec{p}_T$, and p_T^{miss} is the magnitude of this vector.

3.3.6 Particle Flow

We mention here that the full sequence from Section 3.3.1 to Section 3.3.5 is called the “particle flow” (or PF) algorithm.

The PF algorithm provides a complete description of every collision, using data from all detectors. It is possible to use only particle flow objects in an analysis; in practice, however, the purity of PF objects is often not good enough for most analyses in the sense that one can dramatically improve background rejection by using quality cuts to select only a subset of PF criteria that meet more stringent requirements than the implicit quality cuts in the PF algorithm. The quality criteria used in this analysis are detailed in Chapter 5.

We note here that while calorimetric objects as found by the PF algorithm are often not used directly in physics analyses, this algorithm still provides an excellent estimate of p_T^{miss} . The reason is that often, a significant number of energy deposits in an event are left unaccounted for due to a failure of the corresponding PF objects to pass the more stringent quality cuts. This is why most analyses that use p_T^{miss} rely on the PF estimate of p_T^{miss} rather than a custom computation of the negative sum of $\vec{p}_T$ over each well-reconstructed particle passing quality criteria: $\vec{p}_T^{\text{miss}}$ as found by the PF algorithm is less susceptible to bias and its uncertainties are better controlled.

Chapter 4

Datasets and Triggers

This Chapter contains details of the datasets used in this analysis – samples of collision data recorded by the detector, as well as simulated samples used to build the background and signal models. In addition, this Chapter also contains details of the HLT selection used for the data and for those simulated samples that are produced with an emulated trigger. We also include details on how the event weight is fine-tuned in order to correct for various sources of error in a direct MC-to-data comparison.

4.1 Samples of Recorded Data

Table 4.1 contains a complete list of Run 2 datasets used in the analysis. Not all recorded data is suitable for physics analyses (for example, if some components of the detector are malfunctioning during data-taking, breaking the assumptions of the PF algorithm). The CMS detector has a team responsible for data certification, which releases JSON files that can be used as filters for physics analyses to only run over that subset of the recorded data taken while the detector was running nominally. We use data from proton-proton collisions taken at $\sqrt{s} = 13$ TeV in 2016, 2017, and 2018. The data is n-tuplized using the widely used `ggNtuplizer` package. The data for 2016, 2017, and 2018 corresponds to a total integrated luminosity of 36.30 fb^{-1} , 41.53 fb^{-1} , and 59.74 fb^{-1} respectively, for a total of 137.6 fb^{-1} .

4.2 Simulated Samples

Monte Carlo (MC) simulation is used to obtain the potential signal yield in the search for the stealth SUSY signal models. Each of the simplified models introduced in Section 2.2.3 is generated with `MADGRAPH5_aMC@NLO` [47, 48] at leading order (LO). `PYTHIA8` is used for particle showering and hadronization. The generator is

Table 4.1: Datasets and JSON masks used in this analysis.

2016 Datasets:	
	/DoubleEG/Run2016B-17Jul2018_ver1-v1/MINIAOD
	/DoubleEG/Run2016B-17Jul2018_ver2-v1/MINIAOD
	/DoubleEG/Run2016C-17Jul2018-v1/MINIAOD
	/DoubleEG/Run2016D-17Jul2018-v1/MINIAOD
	/DoubleEG/Run2016E-17Jul2018-v1/MINIAOD
	/DoubleEG/Run2016F-17Jul2018-v1/MINIAOD
	/DoubleEG/Run2016G-17Jul2018-v1/MINIAOD
	/DoubleEG/Run2016H-17Jul2018-v1/MINIAOD
	/SinglePhoton/Run2016B-17Jul2018_ver2-v1/MINIAOD
	/SinglePhoton/Run2016C-17Jul2018-v1/MINIAOD
	/SinglePhoton/Run2016D-17Jul2018-v1/MINIAOD
	/SinglePhoton/Run2016E-17Jul2018-v1/MINIAOD
	/SinglePhoton/Run2016F-17Jul2018-v1/MINIAOD
	/SinglePhoton/Run2016G-17Jul2018-v1/MINIAOD
	/SinglePhoton/Run2016H-17Jul2018-v1/MINIAOD
2016 JSON: Cert_271036-284044_13TeV_ReReco_07Aug2017_Collisions16_JSON.txt	
2017 Datasets:	
	/DoubleEG/Run2017B-31Mar2018-v1/MINIAOD
	/DoubleEG/Run2017C-31Mar2018-v1/MINIAOD
	/DoubleEG/Run2017D-31Mar2018-v1/MINIAOD
	/DoubleEG/Run2017E-31Mar2018-v1/MINIAOD
	/DoubleEG/Run2017F-31Mar2018-v1/MINIAOD
	/SinglePhoton/Run2017B-31Mar2018-v1/MINIAOD
	/SinglePhoton/Run2017C-31Mar2018-v1/MINIAOD
	/SinglePhoton/Run2017D-31Mar2018-v1/MINIAOD
	/SinglePhoton/Run2017E-31Mar2018-v1/MINIAOD
	/SinglePhoton/Run2017F-31Mar2018-v1/MINIAOD
2017 JSON: Cert_294927-306462_13TeV_EOY2017ReReco_Collisions17_JSON.txt	
2018 Datasets:	
	/EGamma/Run2018A-17Sep2018-v2/MINIAOD
	/EGamma/Run2018B-17Sep2018-v1/MINIAOD
	/EGamma/Run2018C-17Sep2018-v1/MINIAOD
	/EGamma/Run2018D-22Jan2019-v2/MINIAOD
2018 JSON: Cert_314472-325175_13TeV_17SeptEarlyReReco2018ABC_PromptEraD_Collisions18_JSON.txt	

Table 4.2: List of MC datasets utilized in this Thesis.

Signal MC:	
di-gluino: /SNS-T5tgStealth_TuneCP2_13TeV-madgraphMLM-pythia8/RunIIFall17MiniAODv2-PuFall17Fast_94X_mc2017_realistic_v15-v1/MiniAODSIM	
di-squark: /SNS-T6tgStealth_TuneCP2_13TeV-madgraphMLM-pythia8/RunIIFall17MiniAODv2-PuFall17Fast_94X_mc2017_realistic_v15-v1/MiniAODSIM	
DiPhotonJetsBox:	
2016: /DiPhotonJetsBox_MG-80toInf_13TeV-Sherpa/RunIISummer16MiniAODv3-PuWoriond17_94X_mcRun2_asymptotic_v3-v2/MiniAODSIM	
2017: /DiPhotonJetsBox_MG-80toInf_13TeV-Sherpa/RunIIFall17MiniAODv2-Pu2017_12Apr2018_94X_mc2017_realistic_v14-v2/MiniAODSIM	
2018: /DiPhotonJetsBox_MG-80toInf_13TeV-Sherpa/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MiniAODSIM	
GJetsHT:	
2016: /GJets_DR-Op4_HT-*_TuneCUETP8M1_13TeV-madgraphMLM-pythia8/RunIISummer16MiniAODv3-PuWoriond17_qcut19_94X_mcRun2_asymptotic_v3-v2/MiniAODSIM	
2017: /GJets_DR-Op4_HT-*_TuneCP5_13TeV-madgraphMLM-pythia8/RunIIFall17MiniAODv2-Pu2017_12Apr2018_94X_mc2017_realistic_v14-v*/MiniAODSIM	
2018: /GJets_DR-Op4_HT-*_TuneCP5_13TeV-madgraphMLM-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MiniAODSIM	
QCDHT:	
2016: /QCD_HT*_TuneCUETP8M1_13TeV-madgraphMLM-pythia8/RunIISummer16MiniAODv3-PuWoriond17_94X_mcRun2_asymptotic_v3_ext1-v2/MiniAODSIM	
2017: /QCD_HT*_TuneCP5_13TeV-madgraph-pythia8/RunIIFall17MiniAODv2-Pu2017_12Apr2018_94X_mc2017_realistic_v14-v2/MiniAODSIM	
2018: /QCD_HT*_TuneCP5_13TeV-madgraphMLM-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MiniAODSIM	
$H \rightarrow \gamma\gamma$ MC:	
2016: /GluGluHToGG_M125_13TeV_amcatnlOfFXFX-pythia8/RunIISummer16MiniAODv3-PuWoriond17_94X_mcRun2_asymptotic_v3_ext2-v2/MiniAODSIM	
2017: /GluGluHToGG_M125_13TeV_amcatnlOfFXFX-pythia8/RunIIFall17MiniAODv2-Pu2017_12Apr2018_94X_mc2017_realistic_v14-v1/MiniAODSIM	
2018: /GluGluHToGG_M125_TuneCP5_13TeV-amcatnlOfFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MiniAODSIM	

run with the CP2 tune, which uses the NNPDF3.1LO parton distribution functions¹, or PDFs [49]. The cross sections of the signal models are calculated with a next-to-next-to-leading order (NNLO) approximation up to next-to-next-to-leading logarithmic (NNLL) accuracy [50, 51, 52, 53, 7]. The CMS detector response is determined via a fast simulation [54].

With the masses of $\tilde{S}$ and S fixed to 100 GeV and 90 GeV respectively, we consider a two-dimensional scan of models by allowing the gluino (squark) mass, labeled as $m_{\tilde{g}}$ ($m_{\tilde{q}}$), to vary in the range $1000 \text{ GeV} < m_{\tilde{g}} < 2350 \text{ GeV}$ ($850 \text{ GeV} < m_{\tilde{q}} < 2050 \text{ GeV}$) in increments of 50 GeV. The neutralino mass is then allowed to range from 100 GeV (200 GeV) to the gluino (squark) mass in variable increments, with more points simulated near the boundaries, where the acceptance of the analysis is often sensitive to small changes. The branching ratio of the squark to chargino and neutralino is taken to be 50% each in all models.

MC simulation is also used to refine the data-driven method by looking at the sum of three primary contributions to the SM background: diphoton (with both reconstructed photons produced in the initial hard scattering), γ +Jet (with one reconstructed photon produced in the initial hard scattering and one produced during jet fragmentation or being a misidentified jet), and multijet (with both reconstructed photons produced during jet fragmentation or being misidentified jets). In each of the three samples, PYTHIA8 is used for showering and hadronization. All samples are generated separately for 2016, 2017, and 2018 running conditions. The multijet and γ +Jet samples are generated with MADGRAPH5_aMC@NLO at leading order. The 2016 subsamples of both the multijet and γ +Jet samples are generated with the CUETP8M1 tune, which uses the NNPDF3.0LO PDFs [55], while the 2017 and 2018 subsamples of both the multijet and γ +Jet samples are generated with the CP5 tune, which uses the NNPDF3.1LO PDFs. The diphoton sample is generated with SHERPA 2.1.1 [56] at leading order and uses the CT10 v0 PDFs from the LHAPDF sets [57] for all three years. The CMS detector response is modeled with GEANT4 [58] in all samples.

Table 4.2 contains a complete list of MC datasets used in the analysis, all of which are n-tuplized using the ggNtuplizer package.

For Stealth MC signal samples, we run the event selection using the same year-dependent selection criteria for 2016, 2017, and 2018 data on the Stealth MC samples as on the data, to obtain year-dependent Stealth MC selections. The final yield is then taken as the sum of the event yields for each year, weighted by the year-dependent luminosity, as described in Chapter 6.

¹Incoming protons are composite particles that are made up of several constituents called “partons”, which include quarks and gluons. Each parton carries a fraction of the energy and momentum of its parent proton, described stochastically by so-called “parton distribution functions” (PDFs). Collisions are typically described as an interaction between single partons of two colliding protons, with probabilities derived from these parton distribution functions.

We use three different MC subsamples to obtain SM background distributions of the discriminating variables used in this analysis, which we need in order to evaluate the adjustments needed to the data-driven background model. As described later in this Thesis in Section 5.3, our signal selection has two reconstructed photons in the final state. The following subsamples are used to simulate the various processes that contribute to a two-photon final state:

- The `DiPhotonJetsBox` subsample simulates processes with both photons produced in the initial $2 \rightarrow N$ scattering. This has contributions from both Born processes (tree-level) and so-called “box processes” (diagrams with an internal loop). This process constitutes by far the largest background for this analysis.
- The `GJetsHT` subsample simulates processes with one photon produced in the initial $2 \rightarrow N$ scattering. The other reconstructed photon in the event is a jet misidentified as a photon or a photon produced during jet fragmentation.
- The `QCDHT` subsample simulates events in which both photons in the final event are either jets misidentified as photons, or photons produced during jet fragmentation.

To ensure that there is no overlap between the three subsamples, we count the number of “prompt” final state photons `nPromptFSPhotons` (using the flag `isPromptFinalState`):

- For the `DiPhotonJetsBox` subsample, we require `nPromptFSPhotons == 2`
- For the `GJetsHT` subsample, we require `nPromptFSPhotons == 1`
- For the `QCDHT` subsample, we require `nPromptFSPhotons == 0`

These are all year-dependent `FullSim` samples and trigger requirements are applied as part of the event selection criteria.

Finally, we use the `FullSim` $H \rightarrow \gamma\gamma$ MC samples listed in Table 4.2 in order to estimate the trigger efficiency of our selections for the triggers selected for each year, as described in Section 4.3.

4.3 Triggers

The CMS trigger menu was modified between 2016 and 2017 to ensure sustainable trigger rates. We use events that pass the single photon and diphoton triggers described in Table 4.3 (according to the selection, which will be described further in Chapter 5).

For a hypothetical signal event that satisfies the event selection criteria, there is a small probability that the event might not activate the diphoton HLT trigger that we include in the data selection, and therefore might not be recorded in the data. The trigger inefficiency cannot be estimated directly from the `FastSim` signal samples




Figure 4.1: HLT efficiencies binned in p_T of the leading photon for the diphoton triggers, as obtained from the $H \rightarrow \gamma\gamma$ samples for 2016 (top), 2017 (middle), and 2018 (bottom) data taking conditions.

Table 4.3: List of triggers used for preselection in data and background MC

Year and Selection	Trigger
2016 (diphoton)	HLT_Diphoton30PV_18PV_R9Id_AND_IsoCaloId_AND_HE_R9Id_DoublePixelVeto_Mass55
2017 and 2018 (diphoton)	HLT_Diphoton30PV_18PV_R9Id_AND_IsoCaloId_AND_HE_R9Id_PixelVeto_Mass55
2016 (single photon)	HLT_Photon175
2017 and 2018 (single photon)	HLT_Photon200

because the HLT emulation is not run on those samples. We assume that the trigger efficiency in signal diphoton events is close to the trigger efficiency in the FullSim $H \rightarrow \gamma\gamma$ samples, which reaches 95% for photons with $p_T > 200$ GeV in data taking conditions in all three years. The trigger efficiency, binned in p_T of the leading photon, is shown in Fig. 4.1. The effect of this inefficiency is taken into account in the signal yield by adding an event weight to each stealth MC event that is equal to the trigger efficiency in the p_T bin of the leading photon in the event. Finally, since single photon selections are only used for the FullSim MC background samples, we use the emulated trigger information directly from MC.

4.4 Event Weights for Simulated Samples

Each MC event is weighted in all calculations involving the event yield from that particular MC dataset. This event weight is calculated as follows:

$$\text{event weight} = \frac{\sigma \times \mathcal{L} \times \text{gen weight}}{\sum_{\text{events}} \text{gen weight}} \times \mathcal{W}. \quad (4.1)$$

Here, $\mathcal{L}$ is the total integrated luminosity of the data-taking period. The gen-level weight is typically 1 for all events, but for the `DiPhotonJetsBox` subsample, higher order corrections to the cross section are implemented by means of these weights. The cross section σ is taken as the standard LHC SUSY $\tilde{g}\tilde{g}$ or $\tilde{q}\tilde{q}$ production cross section for Stealth MC. Note that the cross section is independent of the neutralino mass $m_{\tilde{\chi}_1^0}$ because the branching fraction for the neutralino decay into singlino and photon is enforced as a parameter in the MC production. For SM MC subsamples (`DiPhotonJetsBox`, `GJetsHT`, and `QCDHT`), the nominal cross sections are obtained from the `GenXSecAnalyzer` tool. Because it is challenging to produce accurate MC samples with many jets, some of which could be reconstructed as photons, we obtain the nominal yield with these values of the cross section, but in addition, we account




Figure 4.2: Distribution of the prefiring probabilities in MC for gluino production (top) and squark production (bottom) samples.




Figure 4.3: Distribution of the applied MC scale factors for gluino production (top) and squark production (bottom) samples.

for potential mismodeling both as a function of the jet multiplicity and as a function of S_T as described in Section 7.2.5. Finally, $\mathcal{W}$ is an additional MC-to-data weight that takes into account the following corrections:

- Because of known issues with ECAL prefiring in 2016 and 2017, there is a probability that an event produced during a collision in these two years that otherwise passes all selection criteria is not recorded in the data because the trigger activates at the wrong time. These prefiring probabilities can be quantified for each event following a standard procedure, and the distribution of these probability values is in Fig. 4.2 for the Stealth SUSY signal samples.
- There is an additional year-dependent catch-all MC-data scale factor to account for potential MC mismodeling for given photon identification criteria. The distribution of these scale factors is in Fig. 4.3 for the Stealth SUSY signal samples.
- Because MC samples are generated assuming an expected pileup distribution that is not the same as the true pileup distribution in the data, MC events are reweighted to account for this difference.
- The HLT trigger is not fully efficient for any of our selections, and this is accounted for as described in Section 4.3. Note that these weights are used only for Stealth MC samples, since the three background MC subsamples are all FullSim samples generated with an emulated trigger.

Uncertainties associated with the prefiring weights, MC scale factors, and HLT inefficiencies are taken into account, as described in Sections 7.3.3, 7.3.4, and 7.3.5 respectively. Potential errors on the pileup reweighting are not expected to affect the analysis.

Chapter 5

Object Definitions and Event Selection

The algorithms introduced briefly in Chapter 3 produce a collection of reconstructed objects for each event, but without a reliable means of mapping each reconstructed object to a physical particle (such as a photon), such objects are not useful in a realistic analysis. In Section 5.1, we describe the various quality criteria used to identify candidate objects as nominal photons or jets to use in the analysis, and the method of reconstructing the missing transverse energy in an event.

In Section 5.2, we define some key variables defined from the properties of objects identified in the event. These variables are used to build selections for the analysis.

The vast majority of events that merely pass the trigger can be ruled out in order to achieve a reasonable signal-to-noise ratio in the search space of this analysis. In Section 5.3 we describe the event-level criteria to filter events into the signal selection or either of the two control selections, and the bins into which these filtered events are distributed.

5.1 Object Definitions and Identification

5.1.1 Definitions of variables used for Particle ID

In the interest of brevity and readability, we have consolidated into this Section the definitions of all quality variables used to filter PF objects for the analysis.

ID variables for Jets

The fraction of energy attributable to neutral hadrons, neutral EM particles, charged hadrons, and charged EM particles is defined as the energy associated with PF objects

of that particular type within the cone defined by the jet, to the total energy of the jet. The charged multiplicity is defined as the number of charged PF objects associated with the jet. The number of constituents is defined as the total number of PF objects associated with the jet.

ID variables for Photons

The variable R_9 is a shower shape variable and is used in the trigger; it is defined as the fraction of the energy deposited in the 3×3 region around the crystal seed to the energy deposited in the 5×5 region around the crystal seed.

$\sigma_{i\eta i\eta}$ is also a shower shape variable, and quantifies the transverse size of the shower along the η direction:

$$\sigma_{i\eta i\eta} := \sqrt{\frac{\sum_{i \in 5 \times 5} w_i (0.0175 \times \eta_i + \eta_{\text{seed}} - \bar{\eta}_{5 \times 5})^2}{\sum_{i \in 5 \times 5} w_i}} \quad (5.1)$$

$$w_i := 4.2 + \log \left(\frac{E_i}{E_{5 \times 5}} \right)$$

where the seed crystal is the crystal that has a local maximum of deposited energy, the 5×5 region corresponds to the set of 5×5 crystals surrounding the seed, and E_i is the energy deposited in crystal i . While $\sigma_{i\eta i\eta}$ is a bit harder to interpret and more expensive to evaluate than R_9 , it performs better on electrons, whose energy deposits tend to spread along ϕ , as already mentioned previously. This is why offline algorithms – algorithms that are run much after the data has been organized and are not time-bound by trigger constraints and so on – use $\sigma_{i\eta i\eta}$ rather than R_9 . As compared to photons and electrons, jets and parton showers typically have much higher values of R_9 and $\sigma_{i\eta i\eta}$.

The variable H/E is defined as the ratio of the energy deposited in the region of HCAL behind an ECAL supercluster, to the energy of the ECAL supercluster. Photons and electrons have much lower values of H/E than jets.

Next, there are various isolation variables: the charged hadron isolation, the neutral hadron isolation, and the photon isolation. These are all defined as the total amount of energy deposited by PF objects of the corresponding type in a cone of size $\Delta R = 0.3$ around the candidate particle (excluding the candidate particle itself in the case of the photon isolation). If there is no kinematic constraint in a physics signature forcing a generated photon or electron to be collinear with jet deposits, then isolation variables are a powerful tool to weed out deposits from constituents of a hadronizing jet that are reconstructed as photons.

All isolation criteria are corrected for pileup effects. First, “rho”, or the average density of pileup per unit area in (η, ϕ) -space, is estimated in the event; next, the “active area” of the PF photon candidate is evaluated following a slightly different clustering algorithm from the usual anti- k_T algorithm [59]. The pileup contribution to the isolation is estimated from these two quantities and subtracted from the isolation variables; this yields the “rho-corrected” isolation variables that are used in photon identification.

5.1.2 Jets

Jets are reconstructed from collections of PF objects via the anti- k_T clustering algorithm with a distance parameter of 0.4. Jet energy corrections are applied to the jets based on the data-taking period. After correction, selected jets are required to have $p_T > 30$ GeV and to be within the range $|\eta| < 2.4$. In addition, jets must have a pileup jet ID discriminant of > 0.61 and must pass the following tight PF jet ID criteria:

- For 2016, 2017, and 2018:
 - Neutral hadron fraction < 0.9
 - Neutral EM fraction < 0.9
 - Number of constituents ≥ 2
 - Charged hadron fraction > 0
 - Charged multiplicity ≥ 1
 - Charged EM fraction < 0.99
- For 2017 and 2018:
 - Neutral hadron fraction < 0.9
 - Neutral EM fraction < 0.9
 - Number of constituents ≥ 2
 - Charged hadron fraction > 0
 - Charged multiplicity ≥ 1

5.1.3 Photons

Photon candidates are reconstructed from energy clusters in the ECAL. As we shall see in Section 5.3, there are three distinct selections in this analysis: the diphoton signal selection, the diphoton control selection, and the single photon selection. To pass either of the two diphoton selections, the leading photon candidate is required to have $p_T > 35$ GeV, and the subleading photon candidate is required to have

$p_T > 25$ GeV. To pass the single photon selection, the photon candidate is required to have $p_T > 200$ GeV, to ensure good trigger efficiency. Photon candidates are required to be reconstructed in the ECAL barrel, with $|\eta| < 1.442$. To reduce the background from electrons faking photons, all photon candidates must also pass a pixel seed veto: photons are vetoed if there exist two or more hits in the pixel detector leading to the ECAL cluster. Finally, all photon candidates are required to pass the following cuts on shower shape and isolation variables, which together define the “medium” ID requirement for photons.

- For 2016:
 - $H/E < 0.0396$
 - $\sigma_{i\eta i\eta} < 0.01022$
 - ρ -corrected PF charged hadron isolation < 0.441
 - ρ -corrected PF neutral hadron isolation $< 2.725 + 0.0148p_T^\gamma + 1.7 \times 10^{-5}(p_T^\gamma)^2$
 - ρ -corrected PF photon isolation $< 2.571 + 0.0047p_T^\gamma$
- For 2017 and 2018:
 - $H/E < 0.02197$
 - $\sigma_{i\eta i\eta} < 0.01015$
 - ρ -corrected PF charged hadron isolation < 1.141
 - ρ -corrected PF neutral hadron isolation $< 1.189 + 0.01512p_T^\gamma + 2.259 \times 10^{-5}(p_T^\gamma)^2$
 - ρ -corrected PF photon isolation $< 2.08 + 0.004017p_T^\gamma$

If a photon passes all kinematic criteria but fails any of the criteria listed here, it is classified as a sub-nominal photon. The primary diphoton selection is built from medium photons alone, while the diphoton control selection includes sub-nominal photons, as described further in Section 5.3.

5.1.4 Missing transverse momentum p_T^{miss}

The missing transverse momentum vector $\vec{p}_T^{\text{miss}}$ is defined as the transverse component of the negative vector sum of the momenta of all reconstructed particles in the event. The missing transverse momentum p_T^{miss} is then the magnitude of this vector. We apply the following additional filters in order to account for known issues in the calculation of p_T^{miss} :

- Primary vertex filter
- Beam halo filter

- HBHE noise filter
- HBHE iso noise filter
- ECAL TP filter
- Bad PF muon filter
- Bad charged hadron filter
- Bad EE SC noise filter
- Global Super-tight halo 2016 filter
- ECAL bad calibration filter

In addition, we correct $\vec{p}_T^{\text{miss}}$ for the increased noise due to radiation-induced degradation of the ECAL endcaps in 2017/2018.

5.2 Key Analysis Variables

5.2.1 Scalar sum of transverse energy: S_T

S_T is defined as the scalar p_T sum of all selected physics objects in the event, as defined above:

$$S_T \equiv \sum_{\text{photons}} |p_T| + \sum_{\text{jets}} |p_T| + p_T^{\text{miss}} \quad (5.2)$$

The sum includes jets, photons, and p_T^{miss} . In order to ensure that photons reconstructed as jets are not double-counted in S_T , we only include, in Eq. 5.2, the p_T of those reconstructed jets that are more than $\Delta R = 0.4$ away from a reconstructed photon.

5.2.2 Jet multiplicity: nJets

The number of jet objects in an event passing the criteria defined in Section 5.1.2 is defined as the jet multiplicity and abbreviated as nJets throughout this Thesis. In order to ensure that photons reconstructed as jets are not counted in the jet multiplicity, we only include jet objects that are more than $\Delta R = 0.4$ away from a reconstructed photon.

5.3 Event Selection Criteria

The object definitions discussed in the preceding Sections are used to construct the signal selection and two control selections based on the number and quality of

reconstructed photons and jets in the event, as well as the scalar sum S_T . An event is added to the diphoton signal selection if, in addition to passing the diphoton HLT triggers in Table 4.3 it has at least two reconstructed jets, $S_T > 1200$ GeV, and two reconstructed photons as defined in the preceding Sections. In addition we also require the invariant mass of the diphoton system to be greater than 90 GeV.




Figure 5.1: Representation of the search regions used in this analysis as a function of S_T and jet multiplicity (outlined in red), as well as the sideband regions used to derive the shape and normalization of the S_T distribution for the background events.

The signal selection is divided into multiple signal bins based on S_T and jet multiplicity, as shown in Fig. 5.1. The low jet multiplicity sideband is made up of events with exactly two jets, and is used to derive the data-driven input to the shape of the S_T distribution at higher jet multiplicities as described in Section 6.2.2. The region $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$ for events with four or more jets is used as a normalization sideband to fix the scale of the S_T shape determined from the 2-jet bin.

Events with four or more jets and $S_T > 1300 \text{ GeV}$ are taken to be the signal region, where stealth SUSY events might appear. This signal region is binned in three jet multiplicities (4 jets, 5 jets, and ≥ 6 jets) and six S_T ranges ($1300 \rightarrow 1400 \text{ GeV}$, $1400 \rightarrow 1500 \text{ GeV}$, $1500 \rightarrow 1700 \text{ GeV}$, $1700 \rightarrow 2000 \text{ GeV}$, $2000 \rightarrow 2500 \text{ GeV}$, and $\geq 2500 \text{ GeV}$) for a total of 18 search bins in this analysis.

In order to estimate uncertainties on mismodeling both S_T and jet multiplicity as discussed in Section 7.2.5, we produce two control selections. Events passing the single photon HLT triggers in Table 4.3, with exactly one reconstructed photon passing the requirements of Section 5.1.3, and with all other requirements the same as the

diphoton signal selection, make up the single photon control selection. Events passing only the diphoton HLT requirement but without two reconstructed photons as defined in Section 5.1.3 make up the diphoton control selection.

To summarize, we have the following event-level cuts:

- (Stealth SUSY MC only, all selections) two true photons with correct (neutralino) parentage
- (Background MC only, all selections) overlap removal, for correct number of prompt final-state photons
- (data + background MC, all selections) diphoton HLT
- (all samples)
 - diphoton signal selection: 2 medium ID photons
 - diphoton control selection: inversion of diphoton signal selection
 - single photon control selection: 1 medium ID photon with high p_T
- (all samples, all selections) Diphoton invariant mass ≥ 90 GeV
- (all samples, all selections) 2 or more jets passing tight PF jet ID
- (all samples, all selections) $S_T > 1200$ GeV (note: while the normalization region starts from 1200 GeV, we have a 200 GeV-wide buffer in order to avoid pathologies near the kernel boundary as described in Section 6.2.1).

Table 5.1: Cut flow efficiencies for the event-level cuts used in the analysis for the nominal selection on data.

Event Cut	E_{global}	E_{N-1}	$E_{\text{cumulative}}$
diphoton HLT	$\frac{257697613}{2057749548} = 12.52\%$	$\frac{1999}{2179} = 91.74\%$	$\frac{257697613}{2057749548} = 12.52\%$
$N_{\text{medium } \gamma} = 2$	$\frac{3437143}{2057749548} = 0.17\%$	$\frac{1999}{3871} = 51.64\%$	$\frac{2499030}{257697613} = 0.97\%$
$m_{\gamma\gamma} \geq 90$ GeV	$\frac{8033979}{2057749548} = 0.39\%$	$\frac{1999}{2279} = 87.71\%$	$\frac{1166394}{2499030} = 46.67\%$
$N_{\text{jets}} \geq 2$	$\frac{841219253}{2057749548} = 40.88\%$	$\frac{1999}{2320} = 86.16\%$	$\frac{115598}{1166394} = 9.91\%$
$S_T \geq 1000$ GeV	$\frac{18028491}{2057749548} = 0.88\%$	$\frac{1999}{115598} = 1.73\%$	$\frac{1999}{115598} = 1.73\%$

Tables 5.1 and 5.2 respectively have the cut flows for these event-level cuts for the data and Stealth MC samples. These tables have three different measures of the cut efficiency, defined as follows:

Table 5.2: Cut flow efficiencies for the event-level cuts used in the analysis for the nominal selection on the di-gluino production (top), di-squark production (bottom).

Event Cut	E_{global}	E_{N-1}	$E_{\text{cumulative}}$
MC gen cuts	$\frac{43441593}{185082921} = 23.47\%$	$\frac{8262669}{9190861} = 89.90\%$	$\frac{43441593}{185082921} = 23.47\%$
$N_{\text{medium } \gamma} = 2$	$\frac{9405500}{185082921} = 5.08\%$	$\frac{8262669}{14850416} = 55.64\%$	$\frac{8420425}{43441593} = 19.38\%$
$m_{\gamma\gamma} \geq 90 \text{ GeV}$	$\frac{18335035}{185082921} = 9.91\%$	$\frac{8262669}{8395928} = 98.41\%$	$\frac{8286942}{8420425} = 98.41\%$
$N_{\text{jets}} \geq 2$	$\frac{184972751}{185082921} = 99.94\%$	$\frac{8262669}{8283630} = 99.75\%$	$\frac{8265348}{8286942} = 99.74\%$
$S_{\text{T}} \geq 1000 \text{ GeV}$	$\frac{183177338}{185082921} = 98.97\%$	$\frac{8262669}{8265348} = 99.97\%$	$\frac{8262669}{8265348} = 99.97\%$
Event Cut	E_{global}	E_{N-1}	$E_{\text{cumulative}}$
MC gen cuts	$\frac{22980069}{106974651} = 21.48\%$	$\frac{6276608}{7011430} = 89.52\%$	$\frac{22980069}{106974651} = 21.48\%$
$N_{\text{medium } \gamma} = 2$	$\frac{7121107}{106974651} = 6.66\%$	$\frac{6276608}{9462898} = 66.33\%$	$\frac{6347238}{22980069} = 27.62\%$
$m_{\gamma\gamma} \geq 90 \text{ GeV}$	$\frac{11254723}{106974651} = 10.52\%$	$\frac{6276608}{6331753} = 99.13\%$	$\frac{6291717}{6347238} = 99.13\%$
$N_{\text{jets}} \geq 2$	$\frac{106799770}{106974651} = 99.84\%$	$\frac{6276608}{6286672} = 99.84\%$	$\frac{6281114}{6291717} = 99.83\%$
$S_{\text{T}} \geq 1000 \text{ GeV}$	$\frac{105024608}{106974651} = 98.18\%$	$\frac{6276608}{6281114} = 99.93\%$	$\frac{6276608}{6281114} = 99.93\%$

$$\begin{aligned}
E_{\text{global}}(\text{cut } i) &\equiv \frac{N(\text{cut } i)}{N_{\text{analyzed}}} \\
E_{N-1}(\text{cut } i) &\equiv \frac{N_{\text{selected}}}{N(\text{cuts } \{1, 2, \dots, (i-1), (i+1), \dots, M\})} \\
E_{\text{cumulative}}(\text{cut } i) &\equiv \frac{N(\text{cuts } \{1, 2, \dots, i\})}{N(\text{cuts } \{1, 2, \dots, (i-1)\})}
\end{aligned} \tag{5.3}$$

where:

- There are M cuts, labeled by integers from 1 to M
- N_{analyzed} is the total number of events analyzed (the total number of events in the diphoton datasets in Table 4.1 for data, and the total number of events in the Signal MC datasets in Table 4.2)
- $N(\text{cut } i)$ is the total number of events passing cut i (regardless of other cuts)
- $\text{cuts } \{i, j, \dots\}$ denotes the intersection of cuts $i, j, \dots$, so that $N(\text{cuts } \{i, j, \dots\})$ is the number of events that pass all of the cuts $i, j, \dots$
- N_{selected} is the final number of selected events after applying all cuts; with this notation, $N_{\text{selected}} \equiv N(\text{cuts } \{1, 2, \dots, M\})$

Chapter 6

Background and Signal Models

This Chapter describes the background and signal models used in this Thesis. Section 6.1 describes S_T scaling, the main idea behind the background modeling. Section 6.2 introduces the full background model used in the analysis. Section 6.3 contains a brief description of the signal model, which is fairly straightforward in this analysis. Section 6.4 describes how the background and signal models are combined to form the full model used for statistical interpretation, and why the data-driven approach makes this full combination somewhat subtle. Finally, Section 6.5 estimates the potential signal contamination in the normalization sidebands used to build the background model, and demonstrates that such potential signal leakage into those sidebands is not a concern for this analysis.

6.1 S_T scaling

To estimate the SM background, we exploit the fact that the distribution of S_T , as defined in Eq. 5.2, is nearly invariant with respect to the number of jets in SM events.

The assumption of S_T scaling was first used in black hole searches [60, 61, 62, 63], and also in the 8 TeV analysis using 2012 data [31]. An illustration of S_T scaling is in Fig. 6.1, which shows normalized S_T distributions at various jet multiplicities in events with no constraint on the number or quality of photons, from the following high- H_T QCD dataset:

```
/QCD_HT*_TuneCP5_13TeV-madgraph-pythia8/  
RunIIFall17MiniA0Dv2-PU2017_12Apr2018_94X_mc2017_realistic_v14-v*/MINIAODSIM
```

It is seen that S_T scaling “turns on” at increasingly higher values of S_T at higher jet multiplicities, but beyond about 1000 GeV, the S_T distributions are all statistically




Figure 6.1: S_T distribution comparisons across various jet multiplicities, for an MC QCD selection with no photon quality criteria, scaled in the normalization range $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$ to the 2-jets distribution. The bottom portion of the plots shows the ratio with respect to the 2-jets distribution.

identical over multiple orders of magnitude.

For this analysis, assuming a gluino (squark) mass of $m_{\tilde{g}}$ ($m_{\tilde{q}}$), the peak of the S_T distribution for signal events lies at approximately $2 \times m_{\tilde{g}}$ ($2 \times m_{\tilde{q}}$); for any $m_{\tilde{q}}$ larger than the limits set in [31], we expect very little SM contribution in the most sensitive bin at large jet multiplicity and S_T , which makes the S_T shape a useful discriminator between signal and background. However, while S_T scaling is expected to hold for a pure QCD background, it may not hold precisely for backgrounds in which there are selection effects from requiring photons over the QCD background. To account for a potential deviation from S_T scaling, we derive adjustments to the scaling assumption by comparing S_T shapes at high and low jet multiplicity in the sum of the three background MC subsamples, as described in the next Sections.

6.2 Background Model

The full background model used in this analysis, $b^{n \text{ jets}}(S_T \text{ bin } i)$, can be expressed as follows:

$$\begin{aligned} b^{n \text{ jets}}(S_T \text{ bin } i) = & \text{Nevts}^{n \text{ jets}}(1200 \text{ GeV} < S_T < 1300 \text{ GeV}) \\ & \times f^{\text{AGK}}(S_T \text{ bin } i) \\ & \times r^{n \text{ jets}}(S_T \text{ bin } i) \end{aligned} \quad (6.1)$$

where $\text{Nevts}^{n \text{ jets}}(1200 \text{ GeV} < S_T < 1300 \text{ GeV})$ is the number of events in the normalization window, used to correctly scale the S_T shape, $f^{\text{AGK}}(S_T \text{ bin } i)$ is obtained from an Adaptive Gaussian Kernel estimated from the data in the low jet multiplicity sideband, as described below in Section 6.2.1, and the multiplicative adjustment $r^{n \text{ jets}}(S_T \text{ bin } i)$ is obtained from the MC background subsamples as described in Section 6.2.2. Values of r away from 1 indicate a deviation from S_T scaling.

6.2.1 Adaptive Gaussian Kernel estimation

The shape of the S_T distribution at nJets=2, the low jet multiplicity sideband, is obtained as an Adaptive Gaussian Kernel (AGK) template, which depends only on one parameter, ρ . Reference [64] is a detailed review of the theoretical background and implementation of this technique. Briefly, the algorithm associates each event in the input dataset with a Gaussian centered at the corresponding S_T value. A first pass, which is a simple sum of these Gaussians, gives the local density of events as a smooth function of S_T ; a second pass modulates the width of the Gaussian associated with each event such that the width is smaller at those values of S_T with a high local density of datapoints (so that small-scale features in the input dataset can be reproduced accurately where the data permits), and larger at those values of S_T with a low local density of datapoints (so that the template is influenced as little as possible by single events at values of S_T with low statistics). The parameter ρ sets a proportionality constant for the modulation; values of ρ much smaller than 1 tend to lead to templates with spurious peaks due to statistical effects from the original dataset, while values of ρ much larger than 1 tend to lead to templates that smooth over potentially useful structure in the data.

Fig. 6.2 shows the data and the AGK template at an optimal value of ρ (determined as explained below in Section 6.2.1) as well as templates with ρ shifted by $\pm 1\sigma$ (where the defining criteria for a 1σ fluctuation are described in Section 7.2.3). Note that lower values of ρ shift the balance towards making the kernel “fit” the data at smaller scales.

In general, in any AGK estimate, as ρ approaches 0, the template reduces to a series of sharp peaks centered at each individual S_T value in the low jet multiplicity




Figure 6.2: AGK Template, obtained with various values of ρ , from the signal selection in observed data at $n\text{Jets} = 2$. The region shaded yellow is the normalization sideband.

sideband in the dataset. Because the background estimate is obtained by integrating the template over the bin range, this approach then becomes equivalent to using a histogram template; at higher values of ρ , the kernel is a smoothed version of the histogram template. Note also that, while a functional fit would be resilient to statistical fluctuations, there is a difficult-to-estimate uncertainty associated with the choice of any particular form for a functional fit. Because the analysis depends crucially on the shape of the S_T distribution, it is desirable to use an approach which is a compromise between these two extremes (histogram template and functional fit); we suggest that the current approach provides a good compromise, and we expect that at some optimal value of ρ , the template is a good approximation to the true shape of the S_T distribution. The current approach also motivates a clean and natural way to estimate the effect of uncertainties on the background modeling on the background yield.

Optimizing ρ

To obtain the optimal value of ρ for this analysis, we evaluate the template from the data at $n\text{Jets}=2$, the $n\text{Jets}$ sideband, at several values of ρ ; at each value of ρ , the negative log likelihood (NLL) provides a measure of how well the template fits the data. Naturally, the NLL of a template with respect to the data used to generate it is not by itself meaningful (operationally, that would lead to a bias towards lower values of ρ when the template is “overfit”); thus, we divide the $n\text{Jets}=2$ data into $N_{\text{folds}} = 3$ equal




Figure 6.3: Goodness-of-fit as a function of ρ , shown here for $N_{\text{folds}} = 3$, with the nominal value of ρ and the $\pm 1\sigma$ fluctuation.

parts. For each value of ρ , the NLL is obtained as the sum of NLLs of each dataset, containing a fraction $\left(\frac{1}{N_{\text{folds}}}\right)$ of the data, with respect to the template obtained from a sum of the remaining $N_{\text{folds}} - 1$ datasets. The value of ρ that minimizes this combined NLL is then taken as the value of ρ with which we obtain the nominal AGK template used in Eq. 6.1.

The result of this exercise is shown in Fig. 6.3. As shown, $\rho_{\text{nominal}} = 1.65$ is the nominal best fit value of ρ , and is the value used for the background model in this analysis. We account for a potential error on the nominal value of ρ as one of the sources of systematic uncertainty, as described in Section 7.2.3.

Note that any data-driven density estimation method has to impose conditions on the template near the minimum value of the template parameter. This gives rise to edge effects because the density below this minimum value is invisible to the algorithm. To mitigate such “edge effects”, we include a 200 GeV-wide buffer below the normalization region; in addition, the data below the buffer is taken to be a reflection of the full data, which is equivalent to the constraint that the slope of the template is zero at the left edge of the buffer.

Expected number of events from S_T AGK template

Once the AGK template $K(S_T)$ is obtained from the signal dataset in the low jet multiplicity sideband, the factor $f^{\text{AGK}}(S_T \text{ bin } i)$ in Eq. 6.1 is obtained from the AGK template scaled to the number of events in the normalization bin $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$:

$$f^{\text{AGK}}(S_T \text{ bin } i) \equiv \frac{\int_{S_T \text{ bin } i} K(S_T) dS_T}{\int_{1200 \text{ GeV} < S_T < 1300 \text{ GeV}} K(S_T) dS_T} \quad (6.2)$$

6.2.2 Adjustments derived from MC

Table 6.1: Various fits considered for the 2-to- n -jets transfer factor.

fit	functional form	Free parameters
Unadjusted 2-jets kernel	$f_{\text{nominal}}(S_T) \times [A]$	A
Unadjusted 2-jets kernel + slope adjustment	$f_{\text{nominal}}(S_T) \times \left[A + m \left(\frac{S_T}{S_{T\text{norm}}} - 1 \right) \right]$	A, m
Unadjusted 2-jets kernel + sqrt adjustment	$f_{\text{nominal}}(S_T) \times \left[A + p \left(\sqrt{\frac{S_T}{S_{T\text{norm}}}} - 1 \right) \right]$	A, p
Unadjusted 2-jets kernel + slope adjustment + sqrt adjustment	$f_{\text{nominal}}(S_T) \times \left[A + m \left(\frac{S_T}{S_{T\text{norm}}} - 1 \right) + p \left(\sqrt{\frac{S_T}{S_{T\text{norm}}}} - 1 \right) \right]$	A, m, p
Unadjusted 2-jets kernel + slope adjustment + sqrt adjustment + quadratic adjustment	$f_{\text{nominal}}(S_T) \times \left[A + m \left(\frac{S_T}{S_{T\text{norm}}} - 1 \right) + p \left(\sqrt{\frac{S_T}{S_{T\text{norm}}}} - 1 \right) + q \left(\left(\frac{S_T}{S_{T\text{norm}}} \right)^2 - 1 \right) \right]$	A, m, p, q

Here, $f_{\text{nominal}}(S_T) \equiv \text{Nevts}^{n \text{ jets}}(1200 \text{ GeV} < S_T < 1300 \text{ GeV}) \times \frac{K^{\text{combined MC}}(S_T)}{\int_{1200 \text{ GeV} < S_T < 1300 \text{ GeV}} K^{\text{combined MC}}(S_T) dS_T}$

where $K^{\text{combined MC}}(S_T)$ is the AGK template found from the normalization sideband in the combined MC background selection.

The adjustments required to the background model due to a potential deviation from S_T scaling are derived from the S_T shapes at various jet multiplicities in the combined (diphoton + (γ +Jet) + multijet, 2016 + 2017 + 2018) MC simulated data. First, we obtain the AGK template from the combined MC at nJets=2, the low jet multiplicity sideband, using the nominal value of ρ obtained from the data as described in Section 6.2.1. The rest of the combined MC simulated data is then fit separately at each jet multiplicity to this AGK template times an adjustment function, with the constants in the function as the parameters of the fit. The fits considered are

listed in Table 6.1, and Section 6.2.2 below details the reason for not including any higher-order corrections beyond a linear term.

Choosing a Functional Form for the Adjustment Term

Table 6.2: p -values corresponding to the observed f -statistic in the combined MC background selection

p-values (f-statistic)	unadjusted vs linear	linear vs (linear+sqrt)	unadjusted vs sqrt	sqrt vs (linear+sqrt)	(linear+sqrt) vs (linear+sqrt +quad)
nJets = 3	0.960	0.839	0.974	0.759	0.896
nJets = 4	1.000	0.971	1.000	0.872	0.863
nJets = 5	0.999	0.088	0.999	0.388	0.102
nJets $\geq$ 6	1.000	0.765	1.000	0.363	0.361

Consider two different functional fits to the same underlying data: a “nominal” fit, and a “complex” fit that is a superset of the nominal fit and has a larger number of parameters. The complex fit will certainly be a better fit to the data, but it is a common question in statistics whether the improvement in the fit is statistically significant. A statistical test called an F -test can be used to answer this question.

In an application of the F -test, we begin by determining m and n , the degrees of freedom for the nominal and complex fits respectively, evaluated as $N^{\text{datapoints}} - N^{\text{parameters}}$. Because the nominal fit has fewer parameters, we have $m > n$. Suppose furthermore that the values of χ^2 are found to be χ_1^2 for the nominal fit and χ_2^2 for the complex fit. Naturally, here $\chi_1^2 > \chi_2^2$. Now consider the f -statistic, which is defined as follows:

$$F \equiv \frac{\frac{\chi_1^2 - \chi_2^2}{m - n}}{\frac{\chi_2^2}{n}} \quad (6.3)$$

If the “null hypothesis” is valid – that is, the data has a distribution compatible with the model defined by the nominal fit – then the F -statistic has an f -distribution with parameters $(m - n, n)$. We can use this distribution to convert the observed f -statistic into a p -value, and these p -values are given in Table 6.2.

Here, we compare the “unadjusted” fit to both a linear fit and a fit with a sqrt term (from Table 6.1). In addition, we compare both the linear and sqrt fits to a combined (linear+sqrt) fit. Finally we also add a quadratic term to the (linear+sqrt) fit and compare the combined (linear+sqrt+quad) to a combined (linear+sqrt) fit.

In Table 6.2, a p -value of x signifies that $(1 - x)$ is the probability of the data being consistent with the simpler hypothesis. In the literature, a typical p -value threshold, used to decide whether to include more terms in a given functional fit, is 0.95; p -values above this threshold typically indicate that it is necessary to add more terms to a nominal fit.

From the first column in Table 6.2, we see that the linear fit seems to be consistently better than the “unadjusted” baseline (as expected). The second and last columns show that additional terms are not necessary in the fit. We also try fits with a sqrt term, and there is no clear evidence that they offer any additional power over the unadjusted baseline than the simpler linear correction. These observations furnish a justification for the choice of the analysis to use a linear adjustment term to quantify adjustments to S_T -scaling in combined MC events.

Best Fits from combined MC

Table 6.3: Best-fit values of A and m with errors in the double photon selections in MC diphoton+ $(\gamma$ +Jet)+multijet events

Best-fit values	A	m
nJets = 3	1.047 ± 0.021	0.255 ± 0.066
nJets = 4	1.036 ± 0.026	0.747 ± 0.089
nJets = 5	0.992 ± 0.037	1.295 ± 0.145
nJets $\geq$ 6	1.110 ± 0.056	2.421 ± 0.226

The AGK template and all nominal fits for higher jet multiplicities are shown in Fig. 6.4. The linear fit is chosen as the nominal fit for this analysis as discussed in Section 6.2.2 and is highlighted in yellow in Table 6.1. For this fit, 100 toy MC fluctuations corresponding to eigenvariations of the fit are shown as translucent lines. These eigenvariations account for the limited statistics in the combined MC selections and are described further in Section 7.2.4. Fig. 6.4 also shows that fits with additional terms beyond the linear adjustment typically fall within the envelope defined by the eigenfluctuations of the linear fit.

This procedure introduces two types of systematic uncertainties. Firstly, there exists an uncertainty due to variations in the fit function resulting from limited statistics in the combined MC selections. Our approach to deal with this uncertainty is described in Section 7.2.4. Secondly, this procedure assumes that the combined MC is an accurate estimate of the true background distributions of jet multiplicity as well as S_T . Both these assumptions are tested and corresponding uncertainties are accounted for separately as described in Section 7.2.5.




Figure 6.4: MC diphoton+(γ +Jet)+multijet data and fits
 (top) AGK template obtained from the signal MC diphoton+(γ +Jet)+multijet selection at 2 jets, the low jet multiplicity sideband.
 (middle) (left: 4 jets, middle: 5 jets, right: ≥ 6 jets) Data at higher jet multiplicities, with all fits overlaid.
 (bottom) (left: 4 jets, middle: 5 jets, right: ≥ 6 jets) Ratio of the data at n jets, as well as each of the fit functions, to the unadjusted normalized AGK template.

Finding adjustments from S_T AGK template

Once the best fit values are found from the combined MC for the adjustments to S_T scaling as a function of jet multiplicity, the factor $r^{n \text{ jets}}(S_T \text{ bin } i)$ in Eq. 6.1 is obtained as the ratio of the adjusted to unadjusted background prediction:

$$r^{n \text{ jets}}(S_T \text{ bin } i) \equiv \frac{\int_{S_T \text{ bin } i} K^{\text{combined MC}}(S_T) \times (\text{adjustment}(S_T, n \text{ jets})) dS_T}{\int_{S_T \text{ bin } i} K^{\text{combined MC}}(S_T) dS_T} \quad (6.4)$$

where $K^{\text{combined MC}}(S_T)$ is the AGK template derived from combined MC events with $n_{\text{Jets}}=2$, and

$$\text{adjustment}(S_T, n \text{ jets}) \equiv A_{\text{best fit}}^{n \text{ jets}} + m_{\text{best fit}}^{n \text{ jets}} \left(\frac{S_T}{S_T^{\text{norm}}} - 1 \right) \quad (6.5)$$

where $A_{\text{best fit}}^{n \text{ jets}}$ and $m_{\text{best fit}}^{n \text{ jets}}$ are obtained using the procedure described in the preceding Section. The best-fit values for $A_{\text{best fit}}^{n \text{ jets}}$ and $p_{\text{best fit}}^{n \text{ jets}}$ can be found in Table 6.3. Note that here, the errors on these parameters are correlated.

6.3 Signal Model

The signal yield in each signal bin throughout the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ or $(m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ parameter-space is obtained by running the signal selection on the Stealth SUSY MC.

The signal yield at any mass point in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ or $(m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ is the sum of the weighted number of events passing the selection out of the events generated at that mass point, with weights determined by Eq. 4.1.

6.4 Full Model

In general, the toolkits used for statistical inference, and interpretations of signal exclusions, rely on one particular convention for combining the background and signal yields in any bin. To be more precise, we assume that in a given bin, if the background yield is b and the signal yield is s , then the full yield at a given signal strength μ is $b + \mu \times s$. This is indeed the starting point for our discussion on statistical analysis in Sections 8.1 and 8.2.

However, data-driven background models, such as the one used in this analysis, are susceptible to potential signal contamination in the data used to derive the background

model. In general, the background prediction is also a function of the signal strength. This means that some changes need to be made to the analysis to be able to use the machinery of Section 8.2 for statistical inference. To start with, Eq. 8.1 has to be modified to make the background model prediction dependent on μ :

$$(S + B)(n \text{ Jets}, S_T \text{ bin } i, \mu) = B(n \text{ Jets}, S_T \text{ bin } i, \mu) + \mu \times S(n \text{ Jets}, S_T \text{ bin } i) \quad (6.6)$$

A dependence on nuisance parameters θ can be introduced in each term above without changing any of the following logic. Therefore symbols for nuisance parameters have been suppressed for clarity.

For the purposes of this analysis, the dependence of B on μ is approximated by noting that the background model for this analysis has the following effective formula in the case $\mu = 0$:

$$\begin{aligned} B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0) &= N(n \text{ Jets}, 1200 \text{ GeV} < S_T < 1300 \text{ GeV}) \\ &\times \frac{N(2 \text{ Jets}, S_T \text{ bin } i)}{N(2 \text{ Jets}, 1200 \text{ GeV} < S_T < 1300 \text{ GeV})} \end{aligned} \quad (6.7)$$

where $N(n \text{ Jets}, S_T \text{ bin } i)$ is the observed number of events in the S_T bin indexed i and at jet multiplicity n Jets. Introducing the shorthand $N(n, i) = N(n \text{ Jets}, S_T \text{ bin } i)$, and with the convention that the index $i = 1$ is the S_T normalization bin, we can write Eq. 6.7 as follows:

$$B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0) = N(n, 1) \times \frac{N(2, i)}{N(2, 1)} \quad (6.8)$$

At nonzero signal strengths, Eq. 6.8 has to be rewritten:

$$B(n \text{ Jets}, S_T \text{ bin } i, \mu) = (N(n, 1) - \mu S(n, 1)) \times \frac{(N(2, i) - \mu S(2, i))}{(N(2, 1) - \mu S(2, 1))} \quad (6.9)$$

where $S(n, i)$ is the expected signal in the corresponding (jet multiplicity, S_T) bin at signal strength 1. Eq. 6.9 can be approximated as follows:

$$\begin{aligned} B(n \text{ Jets}, S_T \text{ bin } i, \mu) &= B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0)^{\text{uncorrected}} \\ &\times [1 + \mu T_1 + \mu^2 T_2 + \mathcal{O}(\mu^3)] \end{aligned} \quad (6.10)$$

where

$$\begin{aligned}
B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0)^{\text{uncorrected}} &= N(n, 1) \times \frac{N(2, i)}{N(2, 1)} \\
T_1 &= \left(-\frac{S(2, i)}{N(2, i)} - \frac{S(n, 1)}{N(n, 1)} + \frac{S(2, 1)}{N(2, 1)} \right) \\
T_2 &= \left[\left(\frac{S(2, 1)}{N(2, 1)} \right)^2 + \frac{S(n, 1)}{N(n, 1)} \times \frac{S(2, i)}{N(2, i)} \right. \\
&\quad \left. - \frac{S(n, 1)}{N(n, 1)} \times \frac{S(2, 1)}{N(2, 1)} - \frac{S(2, i)}{N(2, i)} \times \frac{S(2, 1)}{N(2, 1)} \right]
\end{aligned} \tag{6.11}$$

Note that if the signal contamination terms $\frac{S(n,1)}{N(n,1)}$ and $\frac{S(2,1)}{N(2,1)}$ are small, then the terms second order and higher can be neglected compared to the linear term (at μ values within exclusion limits). Thus, the full corrected background model can be written as:

$$\begin{aligned}
(S + B)(n \text{ Jets}, S_T \text{ bin } i, \mu) &\approx B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0)^{\text{uncorrected}} \times (1 + \mu T_1) \\
&\quad + \mu \times S(n \text{ Jets}, S_T \text{ bin } i) \\
&= \left(B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0)^{\text{uncorrected}} \right) + \mu \times \\
&\quad \left(S(n \text{ Jets}, S_T \text{ bin } i) \right. \\
&\quad \left. + T_1 \times B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0)^{\text{uncorrected}} \right)
\end{aligned} \tag{6.12}$$

Therefore, accounting for signal contamination is equivalent to introducing a shift in the signal prediction:

$$S(n \text{ Jets}, S_T \text{ bin } i) \rightarrow S(n \text{ Jets}, S_T \text{ bin } i) + T_1 \times B(n \text{ Jets}, S_T \text{ bin } i, \mu = 0)^{\text{uncorrected}} \tag{6.13}$$

where the term T_1 is defined in Eq. 6.11. To ensure that the approximations used here are valid, we restrict our inference to those mass bins in which the terms $\frac{S(n,1)}{N(n,1)}$ and $\frac{S(2,1)}{N(2,1)}$ are below 0.1. This rules out a small region of the $(m_{\tilde{g}} \text{ or } m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ parameter-space: some parameter-space bins at very low gluino/squark masses, and

some parameter-space bins in which the neutralino mass is within about 50 GeV of the gluino/squark mass.

Implementing the shift of Eq. 6.13 allows us to directly use the usual CMS libraries for statistical inference, because the full model with this shift is now equivalent to Eq. 8.1. This also allows us to interpret the upper limits on production cross section similarly to all other CMS analyses.

6.5 Potential Signal Contamination

Throughout the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ or $(m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ parameter-space, we obtain an estimate of the potential signal contamination in the both sidebands of this analysis – the normalization sideband $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$ at all nJets multiplicities, as well all S_T bins in the nJets sideband nJets=2 – as the fraction of the weighted number of events obtained from the MC sample to the observed number of events. For the nJets sideband nJets=2, these plots are shown in Figs. 6.5 and 6.6. For the normalization sideband $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$, these plots are shown in Fig. 6.7. Note that for inference we only use those mass points in the parameter-space in which the potential signal contamination in all S_T bins is below 10%.

In Section 6.4, we derived the correction to the signal required to account for the signal contamination. In Figs. 6.8 and 6.11, we show the ratio of the term $B^{\text{uncorrected}} \times T_1$ to the term $S + B^{\text{uncorrected}}$ in Eq. 6.12 for each signal bin. In other words, we plot here the fractional correction to the signal model at signal strength $\mu = 1$.

These plots all show that potential signal contamination is small enough not to contribute to errors in the statistical inference.




Figure 6.5: Potential signal contamination for di-gluino events with 2 jets in all S_T bins.




Figure 6.6: Potential signal contamination for di-squark events with 2 jets in all S_T bins.




Figure 6.7: Potential signal contamination in the normalization range $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$ for di-gluino events (top) and di-squark events (bottom) with 4, 5, and ≥ 6 jets.




Figure 6.8: Ratio of signal correction to expected signal for di-gluino events in all signal bins, $S_T < 1700$ GeV.




Figure 6.9: Ratio of signal correction to expected signal for di-gluino events in all signal bins, $S_T > 1700$ GeV.




Figure 6.10: Ratio of signal correction to expected signal for di-squark events in all signal bins, $S_T < 1700$ GeV.




Figure 6.11: Ratio of signal correction to expected signal for di-squark events in all signal bins, $S_T > 1700$ GeV.

Chapter 7

Uncertainties

7.1 Overview of Uncertainties

In Chapter 6, we described how we construct the background and signal models for this analysis. However, in the interests of notational clarity, Eq. 6.1, Eq. 6.6 (and the surrounding discussions in that Chapter) somewhat obscure the fact that all background and signal predictions are subject to uncertainties from several sources, and that in many cases the effect of such uncertainties could potentially be that the true background or signal yield in a given bin is different from its nominal prediction by an amount that is comparable to the yield itself. In the following Sections, we consider several sources of uncertainty and describe in detail how we account for each. The remainder of this Section is a brief overview of all uncertainties.

The background prediction is obtained from Eq. 6.1; in this prediction, there is a statistical uncertainty on the factor $\text{Nevts}^{n \text{ jets}}$ ($1200 \text{ GeV} < S_T < 1300 \text{ GeV}$). The factors $f^{\text{AGK}}(S_T \text{ bin } i)$ are affected by two independent sources of uncertainty: an uncertainty on the template shape due to the limited statistics at $n\text{Jets}=2$ (the low jet multiplicity sideband, from which the factors $f^{\text{AGK}}(S_T \text{ bin } i)$ are derived), and an uncertainty on the parameter ρ . The factors $r^{n \text{ jets}}(S_T \text{ bin } i)$ are affected by three independent sources of uncertainty: an uncertainty on the fit function (from which the nominal adjustment is evaluated) due to limited statistics in the combined MC selection, and uncertainties due to potential mismodeling, in the combined MC selection, of both the jet multiplicity and S_T . All these uncertainties are described further in Section 7.2.

For the signal prediction, firstly, there is a statistical uncertainty on the weighted number of MC events. Next, there is an uncertainty on the luminosity, which manifests as a global uncertainty in the signal yield. There is also an uncertainty on the MC due to a systematic overcorrection or undercorrection of the jet energy scale (which

Table 7.1: Overview of uncertainties with approximate values.

Source of uncertainty	Depends on ...	Estimated fractional uncertainty
(data) statistics in normalization bin	jet multiplicity	15% $\rightarrow$ 22%
(data) Template shape	S_T bin	10% $\rightarrow$ 25%
(data) Choice of ρ	S_T bin	$< 3\%$
(data) statistics in MC-derived adjustment	S_T bin, jet multiplicity	$< 7\%$ (leading mode), $< 3\%$ (subleading mode)
(data) jet multiplicity mismodeling	S_T bin, jet multiplicity	0% $\rightarrow$ 8%
(data) S_T mismodeling	S_T bin, jet multiplicity	10%
(MC) Statistics	$(m_{\tilde{g}}/m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ bin, S_T bin, jet multiplicity	$\sim 1\% \rightarrow 5\%$
(MC) Jet Energy Corrections	$(m_{\tilde{g}}/m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ bin, S_T bin, jet multiplicity	$\sim 5\% \rightarrow 15\%$
(MC) Unclustered MET	$(m_{\tilde{g}}/m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ bin, S_T bin, jet multiplicity	$\sim 0\% \rightarrow 5\%$
(MC) Jet Energy Resolution	$(m_{\tilde{g}}/m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ bin, S_T bin, jet multiplicity	$\sim < 0.1\%$
(MC) Prefiring weights	$(m_{\tilde{g}}/m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ bin, S_T bin, jet multiplicity	$\sim 5\% \rightarrow 10\%$
(MC) Photon scale factor	$(m_{\tilde{g}}/m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ bin, S_T bin, jet multiplicity	$\sim 5\% \rightarrow 10\%$
(MC) Luminosity	constant	1.6%
(MC) HLT inefficiency	S_T bin, jet multiplicity	2%

For data uncertainties, the values here are from the the highest S_T bin (which is most sensitive at high $m_{\tilde{g}}/m_{\tilde{q}}$). For all MC uncertainties, the values cited are typical in the region of the $(m_{\tilde{g}}/m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ parameter-space close to the expected limit contours, and in the sensitive (S_T , jet multiplicity) bins. Note that these estimates do not account for bin-by-bin correlations, so it is possible for a nuisance parameter to have a smaller or larger impact on the signal yield than one would expect from this table alone.

shifts both the jet multiplicity and S_T distributions). Jet energy resolution effects and unclustered energy corrections also tend to impact p_T^{miss} and by extension the S_T distribution; we include these as sources of uncertainty on the MC even though we expect them to be small because p_T^{miss} is typically a small fraction of overall S_T . We use prefiring event weights on the MC samples to correct for ECAL prefiring, which is a known issue with 2016 and 2017 data; the errors on these weights contribute to the MC uncertainty. Next, it is standard practice to use a catch-all photon scale factor to account for general discrepancies between data and MC photon distributions, where events are weighted depending on the kinematic properties of photons in the event. We include uncertainties on these scale factors in the MC systematics. Finally, there is an error in the measurement of HLT efficiency due to limited statistics in the $H \rightarrow \gamma\gamma$ MC from which these efficiencies are derived. Table 7.1 contains estimates of these systematic effects. All these uncertainties are described further in Section 7.3.

Table 7.1 provides a general overview of all uncertainties considered in this analysis, with broad estimates of the impact of each uncertainty on the signal or background yield.

7.2 Uncertainties on Background Model

7.2.1 Statistical uncertainty on number of events in the normalization bin

The fractional uncertainty on the number of events in the normalization bin is obtained from the Poisson error on the number of events in the normalization region in each jet multiplicity bin. This uncertainty is implemented as a set of nuisance parameters, correlated across all S_T bins, with one independent nuisance parameter for each jet multiplicity bin. Each nuisance parameter is a multiplicative correction to the background prediction.

7.2.2 Systematic uncertainty due to fluctuations in template shape

Apart from the nominal AGK template used for the background prediction, there is a range of templates that are statistically compatible with the data at nJets=2 (the low jet multiplicity sideband). To obtain an estimate of the uncertainty on the template due to statistical fluctuations in the data, we use the nominal template and generate 1000 toy MC datasets, shown in Fig. 7.1; by construction, each dataset is statistically compatible with the input data used to generate the nominal template. For each dataset, we obtain its own independent AGK estimate and generate a histogram of $\Delta(S_T \text{ bin } i)$, where:




Figure 7.1: The top plot shows all toy MC datasets for the signal selection (in black), with individual templates (in blue) obtained from those datasets. Original template and data overlaid in red, and the normalization region is shown in yellow. The six plots on the bottom are the individual $\Delta(S_T \text{ bin } i)$ distributions in each signal S_T bin.

$$\Delta(S_T \text{ bin } i) \equiv \left(\frac{f^{\text{AGK, toy MC}}(S_T \text{ bin } i)}{f^{\text{AGK}}(S_T \text{ bin } i)} \right) - 1 \quad (7.1)$$

For each S_T bin, this generates a histogram with 1000 entries, also shown in Fig. 7.1. $\Delta(S_T \text{ bin } i)$ has a distribution centered at 0, and the RMS of this distribution is a measure of the fractional uncertainty on the variation of $f^{\text{AGK}}(S_T \text{ bin } i)$ across templates that are all compatible with the data at nJets=2 (the low jet multiplicity sideband).

This uncertainty is implemented as a set of nuisance parameters, correlated across all jet multiplicity bins, with one independent nuisance parameter for each S_T bin. Each nuisance parameter is a multiplicative correction to the background prediction.

7.2.3 Systematic uncertainty due to ρ

There is also some variation in $f^{\text{AGK}}(S_T \text{ bin } i)$ with a variation in the AGK parameter ρ . To obtain the fractional uncertainty due to this possible variation, we examine the NLL distributions of Fig. 6.3. The $\pm 1\sigma$ fluctuations of ρ are obtained as the values of ρ on either side of ρ_{nominal} that satisfy $\text{NLL}(\rho_{\pm 1\sigma}) = \text{NLL}(\rho_{\text{nominal}}) + 0.5$. These 1σ -shifted templates are shown alongside the nominal templates in Fig. 6.2. In each S_T bin, the average of the fractional deviation from nominal of $f^{\text{AGK}}(S_T \text{ bin } i)$ found from these two templates is taken as a measure of the fractional uncertainty on $f^{\text{AGK}}(S_T \text{ bin } i)$ due to the uncertainty on ρ .

This uncertainty is implemented as a set of nuisance parameters, correlated across all jet multiplicity bins, with one independent nuisance parameter for each S_T bin. Each nuisance parameter is a multiplicative correction to the background prediction.

7.2.4 Systematic uncertainty on MC-derived adjustments

As shown in Fig. 6.4, limited statistics in the nJets sideband lead to an uncertainty on the fit function. In particular, the best fit values of the two parameters A and m are subject to statistical fluctuations. To estimate the systematic uncertainty due to these fluctuations, one approach would be to fluctuate A and m up or down by the uncertainty found by the fitting procedure. But this is not quite correct, because fluctuations along A and fluctuations along m are correlated. For example, an upward fluctuation of A and an upward fluctuation of m would result in an adjustment that significantly enhances the background prediction at, say, $S_T > 2500$ GeV, by an amount that we cannot entirely justify given the statistical fluctuation in the number of events in that bin. The issue is that we are performing a multidimensional fit here. While the fitting procedure correctly assigns a high value of χ^2 to simultaneous upward fluctuations of A and m , this is not correctly accounted for if we treat these fluctuations as independent while estimating their effect on the background model.

To account for this, note that, in the two-dimensional (A, m) parameter-space, we have:

$$\chi^2 = (A \ m) \begin{pmatrix} C_{AA} & C_{Am} \\ C_{mA} & C_{mm} \end{pmatrix} \begin{pmatrix} A \\ m \end{pmatrix} \quad (7.2)$$

where C_{AA} , C_{mm} and $C_{Am} = C_{mA}$ are the elements of the covariance matrix. If the two variables were uncorrelated, then the off-diagonal elements would be 0, and in that case, a procedure that treats fluctuations of those two variables separately would not suffer from the issue described in the preceding paragraph. Thus, instead of looking at fluctuations along A and m , we first find the eigendecomposition of the covariance matrix found by the fitting procedure. The normalized eigenvectors of this covariance matrix furnish a basis in which the covariance matrix is diagonal as required for estimating the effects of systematic fluctuations on background yield.

Here, fluctuations in the (A, m) parameter-space along the two normalized eigenvectors are treated as independent sources of systematic uncertainty. We obtain the $\pm 1\sigma$ fluctuations as follows:

$$\begin{pmatrix} A \\ m \end{pmatrix}_{\pm 1\sigma, \text{ eigenmode } i} = \begin{pmatrix} A \\ m \end{pmatrix}_{\text{nominal}} \pm \sqrt{\lambda_i} \begin{pmatrix} a \\ b \end{pmatrix}_{\text{eigenmode } i} \quad (7.3)$$

where $\sqrt{\lambda_i}$ is the eigenvalue of the covariance matrix corresponding to the normalized eigenvector $\begin{pmatrix} a \\ b \end{pmatrix}$. The fractional deviation of $r^{n \text{ jets}}(S_T \text{ bin } i)$ in these $\pm 1\sigma$ fluctuations as compared to the nominal fit is then taken as the estimate of the systematic uncertainty on MC-derived adjustments. This uncertainty is implemented as a set of nuisance parameters, correlated across all S_T bins, with two independent nuisance parameters (corresponding to the two eigenmodes) for each jet multiplicity bin. Each nuisance parameter is a multiplicative correction to the background prediction.

7.2.5 Systematic uncertainties due to mismodeling

Using the procedure of Section 6.2.2 to obtain adjustments to MC assumes that the MC accurately replicates both the S_T and jet multiplicity distributions in data. However, SM background processes are challenging to simulate accurately in the phase space of this analysis. Note that because the S_T of an event is determined largely at the level of the initial $2 \rightarrow N$ process while the jet multiplicity is determined by jet fragmentation, one can treat mismodeling in S_T independently of mismodeling in jet multiplicity. This Section describes how we estimate uncertainties due to both sources.

First, a mismodeling of the jet multiplicity in any of the background subsamples would manifest as an overall error in the nominal cross section for that subsample, thus affecting the estimated adjustments in Eq. 6.4 and the final background model. More specifically, because we normalize the sum of background subsamples in the range $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$, such mismodeling would influence the relative contribution of any one of the three subsamples to the MC-derived adjustments applied to the data-driven background. In this analysis, to estimate the uncertainty associated with a mismodeling of jet multiplicity, we consider six modified combined MC selections. In each combined sample, for each jet multiplicity, one of the three background subsamples is given an additional uniform MC weight of either 2.0, or 0.5, thus enhancing or suppressing its relative contribution to the full MC combination.




Figure 7.2: Event yield ratios, (shifted total MC) / (nominal total MC), as a function of S_T in the 4 jets bin. The top and bottom rows correspond to the modified combination with the relative contribution of each subsample respectively enhanced and suppressed by a factor of 2. Left: DiPhotonJetsBox shifts, middle: GJetsHT shifts, right: QCDHT shifts.

These six combinations are obtained separately for each jet multiplicity. In short, each combination emulates what the true background might look like if any one of the three contributions were relatively enhanced or suppressed by a factor of 2, and the non-closure of Eq. 6.1 (as applied to the combined sample) is taken as a measure of the uncertainty.

In more detail, the systematic uncertainty associated with jet multiplicity mismodeling is obtained as follows:

1. The “shifted” MC combination is generated by weighting any one of the three subsamples up or down relative to the other two by a factor of 2.
2. The S_T distribution at nJets=2 (the low jet multiplicity sideband) is adjusted




Figure 7.3: Event yield ratios, (shifted total MC) / (nominal total MC), as a function of S_T in the 5 jets bin. The top and bottom rows correspond to the modified combination with the relative contribution of each subsample respectively enhanced and suppressed by a factor of 2. Left: DiPhotonJetsBox shifts, middle: GJetsHT shifts, right: QCDHT shifts.




Figure 7.4: Event yield ratios, (shifted total MC) / (nominal total MC), as a function of S_T in the ≥ 6 jets bin. The top and bottom rows correspond to the modified combination with the relative contribution of each subsample respectively enhanced and suppressed by a factor of 2. Left: DiPhotonJetsBox shifts, middle: GJetsHT shifts, right: QCDHT shifts.

using the values of A and m from Table 6.3 (i.e. from the nominal combined MC selection), to obtain the background prediction at 4, 5, or ≥ 6 jets.

3. The ratio of the actual event yield to the background prediction is plotted in black in Figs. 7.2-7.4. Any potential error on the value of $r^{n \text{ jets}}(S_T \text{ bin } i)$ due to an error on the assumed background composition, caused by mismodeling of the jet multiplicity resulting in an error on the nominal cross section, is expected to show up as a deviation of these ratios from 1.0.
4. To minimize the impact of statistical fluctuations, these ratios are fit to a straight line. The parameters and associated errors of the linear fit are shown in the top panel in Figs. 7.2-7.4, and the best fits are plotted with a solid blue line. The dashed blue lines are obtained by shifting the constant and slope terms by their errors. Also shown is a shaded green region in the background representing the expected effect of the statistical error on the fit parameters A and m , which is already accounted for in Section 7.2.4. It is seen that the degree of jet multiplicity mismodeling is typically smaller than the statistical uncertainty on the adjustment.
5. In a given S_T bin, the maximum deviation from 1.0 of the linear smoothed ratio over the six combinations (the ratio in the combination with the largest absolute value for the slope) is taken as an estimate of the fractional uncertainty due to a mismodeling of the jet multiplicity.

This uncertainty is implemented as a set of uncorrelated nuisance parameters, one for each bin in S_T and jet multiplicity.

Second, the uncertainty due to a potential MC mismodeling in S_T is estimated by considering the normalized S_T shapes in both the control selections. As described in Section 5.3, these two control selections are orthogonal to the signal selection: a single photon control selection with exactly one reconstructed photon, and a double photon control selection with all events that pass the HLT preselection (but do not clear the signal selection).

The approach we use is to perform a closure test in both control selections, as follows:

1. The S_T distributions are obtained for the combined MC selection and data, for both control selections, as a function of jet multiplicity, using the nominal weights. The MC distributions are scaled to the data and plotted in Figs. 7.5 and 7.6 for the single photon and diphoton control samples respectively. Because the potential signal contamination needs to be checked explicitly for the double photon control selection, we also obtain estimates of the MC yield at one representative point in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ parameter-space at $m_{\tilde{g}} = 2000 \text{ GeV}$ and $m_{\tilde{\chi}_1^0} = 1000 \text{ GeV}$, plotted as a dashed line in 7.6, showing that potential signal contamination is not a concern in the double photon control selection. In addition, ratios of the S_T distribution at higher jet multiplicities to the S_T




Figure 7.5: Distributions of S_T in the single photon control selection. The plots on the left are S_T distributions, and the plots on the right have the ratios of the data to the combined MC selection. Black: data. Backgrounds: red $\rightarrow$ DiPhotonJetsBox, blue $\rightarrow$ GJetsHT, green $\rightarrow$ QCDHT




Figure 7.6: Distributions of S_T in the double photon control selection. The plots on the left are S_T distributions, and the plots on the right have the ratios of the data to the combined MC selection. The dotted line represents potential signal contamination from Stealth samples for one representative point in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ parameter-space at $m_{\tilde{g}} = 2000$ GeV and $m_{\tilde{\chi}_1^0} = 1000$ GeV.

Black: data. Backgrounds: red $\rightarrow$ DiPhotonJetsBox, blue $\rightarrow$ GJetsHT, green $\rightarrow$ QCDHT




Figure 7.7: Values of the 2-to- n -jets adjustment in the single photon control selection. Left: data, right: combined MC selection.




Figure 7.8: Values of the 2-to- n -jets adjustment in the double photon control selection. Left: data, right: combined MC selection.




Figure 7.9: Ratio of $\left(\frac{\text{data}}{\text{MC}} \text{ at } n \text{ jets}\right)$ to $\left(\frac{\text{data}}{\text{MC}} \text{ at } 2 \text{ jets}\right)$, for $n = 4$ (left), $n = 5$ (middle), $n = 6$ (right), in the single photon control selection.




Figure 7.10: Ratio of $(\frac{\text{data}}{\text{MC}} \text{ at } n \text{ jets})$ to $(\frac{\text{data}}{\text{MC}} \text{ at } 2 \text{ jets})$, for $n = 4$ (left), $n = 5$ (middle), $n = 6$ (right), in the double photon control selection.

distribution at 2 jets for both data and the combined background MC are shown for both control selections in Figs. 7.7 and 7.8.

2. Note that if there is any mismodeling in MC as a function of S_T , its effect on the background prediction would be to change the data/MC ratio at high jet multiplicity relative to the data/MC ratio at $n_{\text{Jets}}=2$ (the low jet multiplicity sideband). Therefore, consider the ratio of the data/MC ratio at high jet multiplicity to the data/MC ratio at $n_{\text{Jets}}=2$ (the low jet multiplicity sideband), which can alternately be read as the ratio of the 2-to- n -jet adjustments observed in data to the 2-to- n -jet adjustments derived from the combined MC selection. This “ratio of ratios” provides a closure test, and in the absence of mismodeling, is expected to be around 1.0 at all values of S_T . The observed values of these “closure ratios” are shown in Figs. 7.9 and 7.10 for the single photon and diphoton control samples respectively.
3. Note that an uncertainty of 10% would be more than adequate to cover this particular source of mismodeling in both these selections; this motivates our choice of a uniform 10% systematic at all values of S_T to account for any mismodeling in S_T .

This uncertainty is implemented as a set of uncorrelated nuisance parameters, one for each bin in S_T and jet multiplicity.

7.3 Uncertainties on Signal Model

7.3.1 Statistical uncertainty on weighted number of MC events

The fractional uncertainty on the signal yield is obtained as the fractional uncertainty on the unweighted total number of events passing the selection criteria. This uncertainty is implemented as a set of independent uncorrelated nuisance parameters, one for each $(S_T, \text{jet multiplicity})$ bin. Each nuisance parameter is a multiplicative correction to the signal prediction.

7.3.2 Systematic uncertainties due to jet energy correction, unclustered energy, and jet energy resolution effects

The corrections to the jet energy scale that are implemented for the MC samples could be systematically underestimated or overestimated. This would lead to a shift in both the jet multiplicity and S_T distributions of the signal MC samples. In addition, p_T^{miss} is vulnerable to unclustered energy corrections and mismeasurement of the jet energy due to finite jet resolution; because p_T^{miss} contributes to S_T , these effects lead to a bias in the S_T distributions of the signal MC samples.

The methods used to estimate these three uncertainties all follow the same general approach. In each case we re-run the full event selection for the MCs by producing two new MC samples – with all jet energies shifted up and down by the (jet-dependent) uncertainty on the jet energy correction in the first case, and by shifting p_T^{miss} in each event up and down by the relevant uncertainty in the latter two cases. Then the fractional difference in the unweighted number of events between these samples and the one with no corrections is taken as the fractional uncertainty on the signal yield resulting from each of these three effects. These uncertainties are each implemented as a single nuisance parameter correlated across all (S_T , jet multiplicity) bins and across both selections. Each nuisance parameter is a multiplicative correction to the signal prediction.

7.3.3 Systematic uncertainty due to prefiring weights

Event-based weights are applied to the MC sample to correct for the known prefiring issue with the 2016 and 2017 data. The prefiring probabilities have an uncertainty which is taken into account by calculating the maximum and minimum event weights consistent with the known uncertainties on the prefiring probabilities; the signal yield is then calculated using these maximum and minimum weights in addition to the nominal weights. The deviation of the signal yield from nominal is then taken as the estimate of the systematic uncertainty due to prefiring weights. This uncertainty is implemented as a single nuisance parameter correlated across all (S_T , jet multiplicity) bins and across both selections. This nuisance parameter is a multiplicative correction to the signal prediction.

7.3.4 Systematic uncertainty due to photon MC scale factors

The photon MC scale factor is a catch-all factor that is included to account for miscellaneous differences between photon distributions in MC and data. Each photon is assigned a weight depending on its kinematic properties. To account for an uncertainty on these scale factors, the weight of each photon is shifted up and down by the uncertainty. The signal yield is then calculated using these shifted weights in addition to the nominal weights, and, as in Section 7.3.3, the deviation of the signal yield is

taken as the estimate of the systematic uncertainty due to photon scale factor errors. This uncertainty is implemented as a single nuisance parameter correlated across all (S_T , jet multiplicity) bins and across both selections. This nuisance parameter is a multiplicative correction to the signal prediction.

7.3.5 Uncertainty due HLT inefficiency

To estimate the uncertainty due to HLT inefficiency, we calculate the signal yield using three event weights for each event: the nominal event weight, obtained by using the nominal efficiencies shown in Fig. 4.1, and maximum and minimum event weights, obtained by shifting these efficiencies up and down by their error in the bin corresponding to the leading photon p_T . The fractional difference between these event yields is taken as an estimate of the uncertainty due to HLT inefficiency. This uncertainty is implemented as a set of independent uncorrelated nuisance parameters, one for each (S_T , jet multiplicity) bin. Each nuisance parameter is a multiplicative correction to the signal prediction.

7.3.6 Systematic uncertainty due to luminosity

Luminosity measurements for the data-taking period are subject to several sources of uncertainty. These are condensed into a single global value annually, which is 1.6% for the Run 2 data-taking period. This uncertainty is implemented as a single nuisance parameter correlated across all (S_T , jet multiplicity) bins and across both selections. This nuisance parameter is a multiplicative correction to the signal prediction.

Chapter 8

Statistical Analysis

8.1 General Overview

A search for physics beyond the Standard Model can generally be framed, in statistical terms, as a hypothesis test. The “null hypothesis” is that LHC collisions follow the predictions of the Standard Model of particle physics, without any beyond-Standard-Model (“BSM”) enhancements. The hypothesis under test is that, in additions to collisions produced by Standard Model processes, there are also collisions produced by processes involving the creation of supersymmetric particles that decay through the Stealth mechanism.

To be more precise, we define some independent search regions (called “bins”). In this analysis, events falling into each of the six independent ranges of S_T and with a given jet multiplicity fall into one bin. Thus the total number of bins is $6(S_T \text{ range possibilities}) \times 3(\text{jet multiplicity possibilities}) = 18$. In each bin, the following is known:

- The number of events observed: obtained from the raw data
- The number of events expected from SM processes, also called the B -only (“background-only”) prediction: obtained from a background model
- The number of events expected from signal processes: obtained from a signal model

In this approach, the signal strength μ is a parameter of the full model. The number of events expected in bin i , N_i^{expected} , is expressed as:

$$N_i^{\text{expected}} = B_i + \mu \times S_i \quad (8.1)$$

where B_i is the background-only prediction and S_i is the signal prediction. Note

that μ is a common multiplicative factor applied to the signal prediction in all bins. This allows us to analyze the shape of the signal (its relative abundance in each bin) independently of its overall normalization. By convention, the overall normalization for S_i is set so that a signal strength $\mu = 1$ corresponds to some standard production cross section that is natural in the theory, so that results from diverse analyses can be interpreted and compared more easily.

The technique detailed in this Chapter allows us to infer, with quantifiable confidence, an upper limit on the signal strength μ for a given signal prediction, background-only prediction, and observation. S_i varies throughout the two-dimensional ($m_{\tilde{g}}$ or $m_{\tilde{q}}$, $m_{\tilde{\chi}_1^0}$) parameter-space, and therefore the limit on μ varies throughout the parameter-space. Of particular interest are those regions in the parameter space where a signal strength of $\mu = 1$ is ruled out by the data, because in these regions, the data are inconsistent with a reasonable theory of Stealth SUSY. In other words, theoretical models corresponding to those regions of the parameter-space can be considered experimentally ruled out. The goal of particle physics experiments is to expand these regions as much as possible (or to examine anomalies further if some particular region cannot be ruled out).

8.2 Procedure for Obtaining Upper Limits

In order to obtain the upper limit on signal strength μ , we introduce a likelihood function that quantifies the consistency between the prediction (from either a background-only model or a full signal+background model) in Section 8.2.1. We account for nuisance parameters in Section 8.2.2. Comparing the values of the likelihood function for the background-only fit versus the (signal+background) fit allows us to construct a test statistic as described in Section 8.2.3. Finally, once the test statistic is known as a function of the signal strength μ , the procedure for obtaining the upper limit is described in Section 8.2.4.

8.2.1 Likelihood Function without nuisance parameters

The likelihood function $L_0(\text{data} \mid \mu)$ quantifies the level of agreement between the observed data and the prediction at a given value of the signal strength parameter μ . In particular, the case $\mu = 0$ corresponds to the background-only model. Suppose the background model and signal model give exact predictions for the expected number of events in all bins. Then, the likelihood function $L_0(\text{data} \mid \mu)$ is just a product of Poisson probabilities, as in any counting experiment:

$$L_0(\text{data} \mid \mu) = \prod_{i \in \text{bins}} \frac{(\mu \times S_i + B_i)^{N_i}}{N_i!} \exp(-\mu \times S_i - B_i) \quad (8.2)$$

where B_i , S_i , μ , and N_i are respectively the background model prediction in bin i , the signal model prediction in bin i , the signal strength, and the observed number of events in bin i .

8.2.2 Treatment of Nuisance Parameters

The background and signal models are, in fact, not as precise as required for Eq. 8.2 to be useful. As described in Chapter 7, there are various parameters that affect the prediction from these models, thereby affecting the likelihood. The true values of these so-called “nuisance parameters” (denoted as θ^j , where j varies from 1 to the number of nuisances K) are not known in advance; however, the distributions of θ^j are typically known, in the sense that the impact of the uncertainty on θ^j can be quantified in each bin i as a known fractional spread $\sigma_i^{j,S}$ or $\sigma_i^{j,B}$ in the signal or background model prediction. Here, we take all nuisance parameters to be normally distributed (i.e. drawn from a Gaussian distribution peaked at 0 with standard deviation 1). For example, if an upward $+1\sigma$ fluctuation in θ^j has the following characteristics:

- There is no effect on the background model prediction in any bin
- The signal model prediction in bin i_1 is enhanced by 10%
- The signal model prediction in bin i_2 is diminished by 5%
- There is no effect on the signal model prediction in any bins other than i_1 and i_2

then, for that nuisance parameter, we would have $\sigma_{i_1}^{j,S} = +0.1$, $\sigma_{i_2}^{j,S} = -0.05$, and $\sigma = 0$ for all other signal predictions and for all background predictions.

To simplify the notation, the vector θ below stands for a collection of all nuisance parameters θ^j . In order to account for nuisance parameters, we make the following changes to the likelihood function:

- The signal and background model predictions in a given bin, S_i and B_i , are now functions of the nuisance parameters: $S_i(\theta)$ and $B_i(\theta)$.
- To the likelihood function, a factor is added that punishes excursions of each nuisance parameter away from their nominal values (which are all 0 by convention).

To be more specific:

1. The impact of nuisance parameter θ^j on the signal or background model prediction in a given bin i is modeled by a multiplicative factor $\left(1 + \sigma_i^{j,S}\right)^{\theta^j}$ or $\left(1 + \sigma_i^{j,B}\right)^{\theta^j}$. Naturally, no such correction is necessary if the nuisance parameter does not actually affect the signal or background model prediction in that bin, which is indeed the behavior of the corrections in the limit $\sigma \rightarrow 0$. Note

that, with this convention, at $\theta^j = 0$, the nominal signal or background model prediction in bin i is maintained; at $\theta^j = -1$, the nominal signal or background model prediction in bin i is multiplied by $\left(1 + \sigma_i^{j,S}\right)^{-1}$ or $\left(1 + \sigma_i^{j,B}\right)^{-1}$; and at $\theta^j = +1$, the nominal signal or background model prediction in bin i is multiplied by $\left(1 + \sigma_i^{j,S}\right)$ or $\left(1 + \sigma_i^{j,B}\right)$.

2. The penalty term in the likelihood for excursions of each nuisance parameter away from its nominal value $\theta^j = 0$ takes the form of a Gaussian, $\frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(\theta^j)^2}{2}\right)$, because that is the assumed distribution of the nuisance parameters.

Putting this all together, the full likelihood function with nuisance parameters included, $L(\text{data} \mid \mu, \theta)$, can be expressed as follows:

$$L(\text{data} \mid \mu, \theta) = p(\theta) \times \prod_{i \in \text{bins}} \frac{(\mu \times S_i(\theta) + B_i(\theta))^{N_i}}{N_i!} \exp(-\mu \times S_i(\theta) - B_i(\theta)) \quad (8.3)$$

where

$$p(\theta) = \prod_{j \in \text{nuisances}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(\theta^j)^2}{2}\right) \quad (8.4)$$

$$S_i(\theta) = \widetilde{S}_i \times \prod_{j \in \text{nuisances}} \left(1 + \sigma_i^{j,S}\right)^{\theta^j} \quad (8.5)$$

$$B_i(\theta) = \widetilde{B}_i \times \prod_{j \in \text{nuisances}} \left(1 + \sigma_i^{j,B}\right)^{\theta^j} \quad (8.6)$$

in which $\widetilde{S}_i$ and $\widetilde{B}_i$ are respectively the nominal signal and background model predictions in bin i (that is, the expected signal and background yields if they did not have to be corrected for nuisances).

It is to be noted that this framework is able to naturally account for correlations between bins, because it is possible to simultaneously enhance or diminish the signal or background model prediction in multiple bins. The simplest example is the fractional uncertainty on signal prediction of 0.016 due to a potential error on the luminosity measurement. In this framework, there is a single nuisance parameter θ^{lumi} for this source of uncertainty, and its effect on the signal model prediction is exactly the same for all signal bins: $\sigma_i^{\text{lumi},S} = 0.016 \forall i$; furthermore, since this uncertainty does not affect any background prediction, we have $\sigma_i^{\text{lumi},B} = 0 \forall i$. Note that negative correlations can also be handled in this framework, by simply having opposite signs for the values of σ in the anticorrelated bins.

8.2.3 Test Statistic

So far, the likelihood function provides information regarding the degree to which the data are compatible with the background-only expectation, as well as with the (signal+background) expectation, for different values of μ . But this does not, by itself, indicate the degree to which the full (signal+background) model for a given value of μ is a better fit to the data than the background-only model. In order to compare the two hypotheses, we need to define a test statistic, which is a scalar into which all this likelihood information is consolidated. The general idea is that the test statistic for random data compatible with the background-only hypothesis has a very different distribution than for random data compatible with the full (signal+background) hypothesis. If both distributions are known, then the observed value of the test statistic (and more specifically, where the observed value lies on both distributions) can be used to infer which of the hypotheses is more likely.

The Neyman-Pearson lemma [65] suggests that the ratio of likelihoods is the most powerful discriminator and should be used to construct such a test statistic, but there is an additional subtlety: the likelihood function depends on the nuisance parameters, and they have to be somehow fixed in evaluating this ratio. We therefore use the so-called “profiled likelihood”, in which we fix the nuisance parameters to those specific values that maximize the likelihood under some appropriate constraints. For the sake of interpretational consistency, both the CMS and ATLAS experiments use the so-called LHC test statistic $\tilde{q}_\mu$ to discriminate between the signal+background and background-only hypotheses. For a given signal strength μ , the LHC test statistic is defined as follows:

$$\tilde{q}_\mu = -2 \ln \frac{L(\text{data} \mid \mu, \hat{\theta}_\mu)}{L(\text{data} \mid \hat{\mu}, \hat{\theta}_{\hat{\mu}})}, \quad 0 \leq \hat{\mu} \leq \mu \quad (8.7)$$

Here, the values of the nuisance parameters θ that maximize the likelihood as defined in Eq. 8.3 for the chosen value of the signal strength μ are denoted by $\hat{\theta}_\mu$. If one allows both the signal strength and nuisances to vary in Eq. 8.3, then the values of signal strength and nuisances that jointly maximize the likelihood are denoted respectively as $\hat{\mu}$ and $\hat{\theta}_{\hat{\mu}}$. Here, while finding the global maximum of the likelihood, the signal strength parameter is only allowed to vary in the range $[0, \mu]$. The lower bound ($\hat{\mu} \geq 0$) ensures that the signal is positive. The upper bound ($\hat{\mu} \leq \mu$) is more subtle. Suppose the global maximum without any constraint, $\widehat{\mu^{\text{global}}}$, were such that $\widehat{\mu^{\text{global}}} > \mu$ (which could happen just due to a statistical fluctuation); then, the procedure would rule out low values of μ in the full (signal+background) model. Since our intention is to set one-sided upper limits, the constraint $\hat{\mu} \leq \mu$ ensures that values of μ between 0 and $\widehat{\mu^{\text{global}}}$ are not ruled out by the procedure.

It is important to note that the data can be either observed data, or pseudo-data generated under either of the two hypotheses (background-only or full signal+background). The observed value of the test statistic – that is, the value obtained by plugging in the observed data into Eq. 8.7 – is distinguished with a special superscript and denoted as $\widetilde{q}_\mu^{\text{obs}}$. With this construction, we always have $\widetilde{q}_\mu^{\text{obs}} \geq 0$; high values of $\widetilde{q}_\mu^{\text{obs}}$ indicate that the data are more likely to be compatible with the full (signal+background) model than with the background-only model. The next step is to quantify this intuition by finding the exact expected distribution of $\widetilde{q}_\mu$ for the two hypotheses under test, denoted by $f(\widetilde{q}_\mu \mid \text{background-only})$ for the background-only hypothesis and $f(\widetilde{q}_\mu \mid \text{signal} + \text{background})$ for the full (background+signal) hypothesis.

One intuitive way to generate these distributions is as follows:

1. Generate pseudo-data using the background-only model. The nuisance parameters used to generate the pseudo-data are fixed to their best fits to the observed data, but are allowed to vary while finding the parameters $\widehat{\theta}_\mu$ and $\widehat{\theta}_{\widehat{\mu}}$ in Eq. 8.7. Each pseudo-data will correspond to a distinct value of $\widetilde{q}_\mu$, and repeated pseudo-experiments allow us to build $f(\widetilde{q}_\mu \mid \text{background-only})$.
2. Repeat the procedure of step 1, but generate pseudo-data using the full (signal+background) model for a given value of μ . This allows us to build $f(\widetilde{q}_\mu \mid \text{signal} + \text{background})$.

The issue with this procedure is that it is computationally expensive, and the computational cost grows rapidly with the number of nuisance parameters. Luckily, there is a useful approximation that is valid in the asymptotic regime (large number of events in all bins). Without the constraint $\widehat{\mu} \geq 0$ in Eq. 8.7, Wilks' theorem implies [66] that $\widetilde{q}_\mu$ has half a χ^2 distribution with one degree of freedom under the null hypothesis. The constraint $\widehat{\mu} \geq 0$ makes this statement no longer true. With this constraint, the distribution of $\widetilde{q}_\mu$ is not a χ^2 distribution, but it is nonetheless analytically expressible. The exact formula can be found in [67].

However, it is also shown in [67] that in the asymptotic regime, the effect of this constraint on the actual value of $\widetilde{q}_\mu$ is small, and therefore, the distributions $f(\widetilde{q}_\mu \mid \text{background-only})$ and $f(\widetilde{q}_\mu \mid \text{signal} + \text{background})$ can be approximated following Wilks' theorem, which is the approach used in the present analysis.

8.2.4 Upper Limit Inference

Once the two distributions $f(\widetilde{q}_\mu \mid \text{background-only})$ and $f(\widetilde{q}_\mu \mid \text{signal} + \text{background})$ are evaluated, consider the following p -value:

$$\text{CL}_{s+b}(\mu) = P\left(\tilde{q}_\mu \geq \widetilde{q_\mu^{\text{obs}}} \mid \text{signal} + \text{background}\right) \quad (8.8)$$

$$= \int_{\widetilde{q_\mu^{\text{obs}}}}^{\infty} f(\tilde{q}_\mu \mid \text{signal} + \text{background}) d\tilde{q}_\mu \quad (8.9)$$

This p -value, called $\text{CL}_{s+b}(\mu)$, is a function of the signal strength parameter μ and quantifies the frequentist probability that the test statistic $\tilde{q}_\mu$ would be its observed value $\widetilde{q_\mu^{\text{obs}}}$ or higher, given the truth of the (signal+background) hypothesis. In other words, the value of $\text{CL}_{s+b}(\mu)$ at a given μ indicates that the full (signal+background) model at signal strength μ can be ruled out with a confidence of $(1 - \text{CL}_{s+b}(\mu))$. Since $\text{CL}_{s+b}(\mu)$ is a monotonic function of μ , this function, by itself, can be used to find upper limits at any desired confidence. For a target confidence α , the “ CL_{s+b} ” upper limit, denoted as $\mu_{\text{CL}_{s+b}}^\alpha$, is defined implicitly by the following equation:

$$\text{CL}_{s+b}\left(\mu_{\text{CL}_{s+b}}^\alpha\right) = 1 - \alpha. \quad (8.10)$$

However, while this approach is entirely statistically sound (and indeed common outside experimental particle physics), the limits obtained using the procedure need to be interpreted carefully. In particular, if there is a statistical fluctuation that pushes the true background yield in the observed data much lower than the value assumed in the background model, then even a signal strength $\mu = 0$ might be ruled out by the procedure of Eq. 8.10. Indeed, one can quantify how often this occurs (assuming the null hypothesis): if analyses report upper limits excluded at a confidence level of 95%, then one would expect the value $\mu = 0$ to be ruled out in 5% of all analyses with such upper limits assuming that the data is distributed exactly as predicted by the SM. In other words, 5% of all analyses that quote upper limits would end up ruling out the Standard Model itself even if it happens to be a perfectly good description of Nature and those analyses have accurate estimates of their uncertainties.

This is not really a statistics issue: pollution of results by statistical fluctuations is in fact unavoidable and, from the point of view of statistics, the only point of concern is whether the procedure for evaluating limits at a claimed confidence has the correct coverage. It is nonetheless a good idea to build some resilience within the upper limits framework to such downward fluctuations of the background. The approach used by both the CMS and ATLAS experiments is to define the following ancillary p -value:

$$\text{CL}_b(\mu) = P\left(\tilde{q}_\mu \geq \widetilde{q_\mu^{\text{obs}}} \mid \text{background-only}\right) \quad (8.11)$$

$$= \int_{\widetilde{q_\mu^{\text{obs}}}}^{\infty} f(\tilde{q}_\mu \mid \text{background-only}) d\tilde{q}_\mu \quad (8.12)$$

and then to define a new variable $\text{CL}_s(\mu)$ as the ratio of two p -values:

$$\text{CL}_s(\mu) = \frac{\text{CL}_{s+b}(\mu)}{\text{CL}_b(\mu)}. \quad (8.13)$$

This ratio factorizes out the effect of statistical fluctuations in the background. Now, for a target confidence α , the “ CL_s ” upper limit, denoted as $\mu_{\text{CL}_s}^\alpha$, is defined implicitly by the following equation:

$$\text{CL}_s(\mu_{\text{CL}_s}^\alpha) = 1 - \alpha. \quad (8.14)$$

This approach sacrifices some ease of interpretation but allows consistency in inference between all analyses. The limits can no longer be easily expressed without reference to the uncommon term CL_s , since there is no widely accepted intuitive interpretation or specific terminology for the ratio of two p -values. On the other hand, analyses, by providing limits that are robust against downward fluctuations of the background, retain their usefulness as guides for exploring new directions in the theory. In addition, the loss in ease of interpretation is at least partly ameliorated by the fact that limits obtained using Eq. 8.14 are, by construction, always more conservative than the upper limits constructed using Eq. 8.10, which can indeed be easily expressed in terms of a meaningful p -value.

Chapter 9

Results

9.1 Overview

All content in this Thesis so far has been building up to this Chapter. Chapters 1 and 2 introduced the SM and described the motivation for a search for Stealth SUSY. Chapter 3 described the LHC complex and the CMS detector, the instruments for this search. Chapters 4 and 5 described the datasets, event selection criteria, and discriminating variables used to distinguish potential signal from SM background. Chapter 6 discussed S_T scaling and described how we build a background and signal model for the analysis. Chapter 7 described the various sources of uncertainty and their effects on the signal and background yields. Finally, Chapter 8 described how statistical inference is run using these models and their uncertainties, and the precise definition and interpretation of upper limits in the analysis.

This Chapter reveals the final results of running the full statistical inference. We present both expected and observed limits on di-gluino and di-squark production. The expected and observed limits are obtained similarly – by comparing the signal yield with the expected background in each signal bin from Eq. 6.1 (for expected limits), or with the observed data (for observed limits). All uncertainties described in Chapter 7 are included. We compute 95% confidence level (CL) upper limits on the expected gluino and squark pair production cross sections in the simplified stealth SUSY models described above using the full dataset from 2016, 2017, and 2018. We use the modified frequentist CL_s method [68, 69, 70] with a log-likelihood ratio test statistic comparing the SM-only hypothesis to the hypothesis that there is an additional contribution from the signal. An asymptotic approximation [67] for the test statistic is used to extract the observed limits. For inference, we only include bins in the $(m_{\tilde{g}} \text{ or } m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ parameter-space in which the full model, as described in Section 6.4, is accurate to within 10%. Specifically, that region of the parameter-space is excluded in which the neutralino mass is within 100 GeV of the gluino (or squark) mass.

Section 9.2 shows the background and signal models for some representative points in the $(m_{\tilde{q}}, m_{\tilde{\chi}_1^0})$ parameter-space, as well as the expected limits, before any fits to the observed data. Section 9.3 shows the observed data, as well as the background and signal models with nuisance parameters fit to the observed data. Section 9.3 also shows the impact plots for the nuisance parameters. Section 9.4 shows the observed limits, the main result of this analysis. Finally, Section 9.5 presents some concluding remarks.

9.2 Expected Distributions and Expected Limits

9.2.1 Expected S_T shapes

Fig. 9.1 shows the expected background S_T distributions, obtained using Eq. 6.1 and setting all nuisance parameters to 0. In addition, a few S_T distributions from selected points in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ parameter-space are also shown. Yellow bands represent the expected errors on the background prediction and are obtained by generating toy MC datasets by varying the nuisance parameters by their uncertainties. Note that this does take bin-by-bin correlations in the nuisance parameters into account. Systematic uncertainties on the signal prediction are not represented in these plots.

9.2.2 Expected Limits

The expected limits obtained as discussed in Section 9.1 are shown in Fig. 9.2 and 9.3 for gluino and squark production respectively.

9.3 Fits to Observed Data

9.3.1 Observed S_T shapes

Fig. 9.4 shows the post-fit expected background S_T distributions in the signal bins, obtained using Eq. 6.1 with the nuisance parameters set to their best fit values at zero signal strength (as opposed to the pre-fit distributions in Fig. 9.1, which were obtained by setting the nuisances to zero). Observations from data are shown as black data-points. Yellow bands represent the uncertainties on the expected background (because these are obtained by sampling nuisances using the covariance matrix before or after the fit, they take bin-by-bin correlations in the nuisance parameters into account). Table 9.1 shows, in tabular form, the same pre- and post-fit predictions and observations, as well as the signal yield for one representative point in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ parameter-space at $m_{\tilde{g}} = 2000$ GeV and $m_{\tilde{\chi}_1^0} = 1000$ GeV. This point is chosen to be close to the expected exclusion contour.




Figure 9.1: Expected pre-fit background S_T distributions. Signal yields are also shown for some representative mass bins. Top: 4 jets bin, middle: 5 jets bin, bottom: ≥ 6 jets bin.




Figure 9.2: The 95% confidence level upper limits on gluino pair production cross sections as a function of gluino and neutralino masses. The red contours show the expected 50% as well as the 16% and 84% exclusions.




Figure 9.3: The 95% confidence level upper limits on squark pair production cross sections as a function of squark and neutralino masses. The red contours show the expected 50% as well as the 16% and 84% exclusions.

Table 9.1: Expected backgrounds and observations at $n\text{Jets} = 4$ (top), $n\text{Jets} = 5$ (middle) and $n\text{Jets} \geq 6$ (bottom). Signal yields are also shown for one specific signal model at $m_{\tilde{g}} = 2000 \text{ GeV}$, $m_{\tilde{\chi}_1^0} = 1000 \text{ GeV}$.

Predicted Background and Observations, 4 Jets				
S_T range [GeV]	background prediction (pre-fit)	example signal model prediction	background prediction (post-fit)	observation
$1300 \leq S_T < 1400$	28.52 ± 5.13	0.00 ± 0.00	37.92 ± 4.00	40
$1400 \leq S_T < 1500$	19.82 ± 3.45	0.00 ± 0.00	28.54 ± 3.35	35
$1500 \leq S_T < 1700$	23.75 ± 4.78	0.00 ± 0.00	30.62 ± 3.17	27
$1700 \leq S_T < 2000$	18.23 ± 4.31	0.00 ± 0.00	26.88 ± 3.22	27
$2000 \leq S_T < 2500$	11.81 ± 2.83	0.03 ± 0.01	16.58 ± 2.66	20
$S_T \geq 2500$	7.82 ± 2.60	0.11 ± 0.02	10.16 ± 1.99	13

Predicted Background and Observations, 5 Jets				
S_T range [GeV]	background prediction (pre-fit)	example signal model prediction	background prediction (post-fit)	observation
$1300 \leq S_T < 1400$	16.51 ± 3.95	0.00 ± 0.00	22.07 ± 2.71	24
$1400 \leq S_T < 1500$	11.91 ± 3.10	0.00 ± 0.00	16.12 ± 2.01	15
$1500 \leq S_T < 1700$	14.93 ± 4.06	0.00 ± 0.00	20.77 ± 2.68	25
$1700 \leq S_T < 2000$	12.17 ± 3.59	0.00 ± 0.00	18.42 ± 2.51	21
$2000 \leq S_T < 2500$	8.43 ± 2.35	0.02 ± 0.01	11.43 ± 2.03	11
$S_T \geq 2500$	6.01 ± 2.37	0.56 ± 0.06	7.54 ± 1.53	7

Predicted Background and Observations, ≥ 6 Jets				
S_T range [GeV]	background prediction (pre-fit)	example signal model prediction	background prediction (post-fit)	observation
$1300 \leq S_T < 1400$	16.80 ± 3.92	0.00 ± 0.00	15.61 ± 2.09	13
$1400 \leq S_T < 1500$	12.73 ± 3.54	0.00 ± 0.00	12.73 ± 1.91	14
$1500 \leq S_T < 1700$	16.82 ± 4.26	0.00 ± 0.00	16.49 ± 2.23	18
$1700 \leq S_T < 2000$	14.62 ± 4.04	0.01 ± 0.00	16.02 ± 2.34	19
$2000 \leq S_T < 2500$	10.78 ± 3.33	0.08 ± 0.02	10.40 ± 1.70	9
$S_T \geq 2500$	8.15 ± 2.78	10.15 ± 0.70	7.07 ± 1.49	4




Figure 9.4: Expected post-fit background S_T distributions and observed data. Signal yield is also shown for some representative mass bins. Top: 4 jets bin, middle: 5 jets bin, bottom: ≥ 6 jets bin.

9.3.2 Nuisance Parameter Impacts




Figure 9.5: Pulls and impacts of nuisance parameters in an unblinded fit to the observed data (1/2).

For any given nuisance parameter, its pull is defined as the amount (in multiples of its uncertainty) by which its best fit value differs from its nominal assumption in the null hypothesis (in our case, 0). Its impact (on the signal strength parameter) is defined as the amount by which the best fit value of the signal strength would change if that particular nuisance were to shift up or down by one unit of its post-fit uncertainty. Impacts and pulls of all nuisance parameters are given in Figs. 9.5 and 9.6. Correlations between the nuisance parameters for the background-only fit and for the signal+background fits are given in Figs. 9.7 and 9.9, respectively, and correlations between the signal bins for the background-only fit and for the signal+background fits are given in Figs. 9.8 and 9.10, respectively.

The nuisances that are significantly pulled for the post-fit prediction are the normalizations for both the 4-jets and 5-jets multiplicities, as shown in Fig. 9.5, which contains




Figure 9.6: Pulls and impacts of nuisance parameters in an unblinded fit to the observed data (2/2).

Correlation matrix of fit parameters




Figure 9.7: Correlation matrix between nuisance parameters for the observed data, background-only fit.

Bin-by-bin correlations




Figure 9.8: Correlation matrix between signal bins for the observed data, background-only fit.

Correlation matrix of fit parameters




Figure 9.9: Correlation matrix between nuisance parameters for the observed data, signal+background fit.

Bin-by-bin correlations




Figure 9.10: Correlation matrix between signal bins for the observed data, signal+background fit.

a complete list of the values by which nuisances are pulled in the post-fit prediction. The general effect of both nuisances is to pull down the expected background at both jet multiplicities.

9.3.3 Validity of the Asymptotic Approximation

Table 9.2: Comparison of the observed limits obtained with the full estimation based on frequentist toys (left) with observed limits obtained with an asymptotic approximation based on Wilks’ theorem (right).

Mass bin	Observed Limit (frequentist MC toys)	Observed Limit (asymptotic approximation)
$m_{\tilde{g}} = 1950, m_{\tilde{\chi}_1^0} = 1850$	0.82 ± 0.06	0.87
$m_{\tilde{g}} = 2100, m_{\tilde{\chi}_1^0} = 1000$	0.75 ± 0.06	0.76
$m_{\tilde{g}} = 1850, m_{\tilde{\chi}_1^0} = 200$	1.15 ± 0.04	1.14

The statistical inference described here relies on the validity of Wilks’ theorem, as discussed in Section 8.2.3. However, this approximation is only provably valid in the asymptotic regime: when the signal and background expectations are large enough. In this analysis, some expected yields are $\mathcal{O}(1)$ in regions of the phase space close to the reported exclusion contours, as seen from Table 9.1. Therefore, we pick three points in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ parameter-space: these three points are chosen to be near the exclusion contours and cover the full range of expected signal distributions. For each of these three points, we obtain full shapes of the profiled likelihood ratio using Monte Carlo (frequentist) toys, and use these shapes to infer the CLs limit without using the asymptotic approximation. We then compare these limits to the ones found using the asymptotic approximation. The results of this analysis are shown in Table 9.2. We see that:

- The observed limits evaluated with the asymptotic approximation are well within the quoted errors in the full calculation
- The differences are much smaller than the $\pm 1\sigma$ contours (both theoretical and experimental)

We conclude that the asymptotic approximation is reasonable, and we find that the greatly increased computational expense of the full $\text{CL}_s(\mu)$ calculation can be avoided.

9.3.4 Analysis of Fits

Most results look nominal. All fits converge. The post-fit nuisance uncertainties are not too much constrained as compared to the pre-fit values. Impacts and pulls look

Table 9.3: Observed number of events and associated Poisson errors in normalization bin and in the whole S_T observation range. Note that the errors presented here are equal to half the Poisson interval for clarity, but in the actual datacards, the up and down Poisson errors are treated correctly as asymmetric errors.

nJets bin	N (norm window)	$\frac{\Delta N}{N}$ (norm window)	N (obs window)	$\frac{\Delta N}{N}$ (obs window)	error ratio (obs/norm)
3 Jets	79	0.119	210	0.071	0.600
4 Jets	49	0.153	162	0.082	0.533
5 Jets	29	0.203	101	0.105	0.514
≥ 6 Jets	30	0.200	77	0.121	0.604

as expected (the nuisances that are expected to have the biggest impact are indeed the ones that do). Correlations in the fits are also generally as expected.

One detail that requires examination is the behavior of the normalization nuisances. These are introduced to account for the statistical fluctuation on the number of events in the S_T normalization bin $1200 \text{ GeV} < S_T < 1300 \text{ GeV}$, and are estimated from the Poisson errors given the observed number of events in the normalization bin.

However, estimating the normalization uncertainty just using the Poisson error on number of events in the normalization bin glosses over the fact that there is also normalization information in all the observation bins, and if the number of events in the whole observation range for S_T is significantly larger than the number of events in the normalization bin, then the post-fit error is indeed expected to be significantly smaller. Table 9.3, which shows the Poisson uncertainties on the number of events in the normalization bin and in the whole observation range of S_T , makes this more concrete.

Note in particular the following observations related to the normalization nuisances:

- The normalization systematic is highly correlated with other systematic effects (as expected). Most of the “significant” correlations in the matrix of correlations between nuisances are related to the normalization.
- As shown in Fig. 9.5, the normalization happens to be pulled upwards in several jet multiplicity bins, which accounts for a large part of the difference between pre- and post-fit backgrounds in Figs. 9.1 and 9.4.

9.4 Observed Limits and Interpretation

The limits obtained as discussed in Section 9.1 from the observed data are shown in Fig. 9.2 and 9.3 for gluino and squark production, respectively.




Figure 9.11: The 95% confidence level upper limits on gluino pair production cross sections as a function of gluino and neutralino masses. The red contours show the expected 50% as well as the 16% and 84% exclusions. The black contours show the observed exclusions assuming NNLO+NNLL signal cross sections, with their 1σ uncertainties.




Figure 9.12: The 95% confidence level upper limits on squark pair production cross sections as a function of squark and neutralino masses. The red contours show the expected 50% as well as the 16% and 84% exclusions. The black contours show the observed exclusions assuming NNLO+NNLL signal cross sections, with their 1σ uncertainties.

For typical neutralino masses, we exclude gluino masses up to 2100 GeV or squark masses up to 1850 GeV at 95% CL. In regions of the parameter-space at low neutralino masses, the neutralino becomes boosted, and the resulting decay products become collimated. This has the effect of causing the photon to fail the isolation criteria, and results in a weakening of the exclusion limits. In regions of the parameter-space where the neutralino mass is very close to the gluino (squark) mass, the jets produced from the gluino decay become soft. This results in a reduction in the overall number of reconstructed jets in the event, and also causes a slight reduction in the excluded gluino masses.

9.5 Concluding Remarks

The analysis detailed in this Thesis represents a $\sim 70\%$ improvement in the reach of the exclusion contour in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ parameter-space as compared to the previous version of the analysis [31], despite significantly relaxed assumptions in the background modeling. These are the most stringent limits to date with these models.

9.5.1 Further Studies and Outlook

The conjecture of supersymmetry has a fairly unique place in LHC phenomenology. In Chapter 1, we saw that the motivation for SUSY is strongest when SUSY is broken at roughly the electroweak scale, and that the motivation becomes weaker at much higher energies. The results from Fig. 2.2, in conjunction with the results from this Thesis and other analyses with unusual signatures, seem to suggest that SUSY is ruled out in a large fraction of this well-motivated parameter-space. Of course, there remains a shrinking fraction of the well-motivated parameter-space for SUSY that still needs to be explored.

Three general ideas used in this analysis appear to have broad applicability beyond the scope of this Thesis:

1. When a background sample is a sum of subsamples with an uncertain mix, the procedure of Section 7.2.5 can be applied in order to estimate uncertainties on the background due to mismodeling.
2. For distributions with a mix of high and low statistics, an adaptive Gaussian kernel estimate helps keep the statistical uncertainties under control. The procedure detailed in Section 6.2.1 can be used to estimate the kernel parameter ρ .
3. In general, a data-driven background modeling procedure has to account for potential signal contamination. If it is not negligible, then the procedure of Section 6.4 can be used to systematically correct the signal to account for such signal contamination.

Extensions of the current analysis are feasible. Two ideas for further studies are covered below; these studies could extend limits on Stealth SUSY and related models.

Stealth SUSY with leptonic final states




Figure 9.13: Diagrams of some potential simplified models with Stealth SUSY not considered in this Thesis. Signatures involve one W and one photon in the final state (top) or two W s in the final state (bottom).

In this Thesis, we have only probed the data for a specific decay chain that involves the decay of gluinos or squarks to two neutralinos, which gives rise to a diphoton final state. However, there are other possible decays compatible with Stealth SUSY, as illustrated in Fig. 9.13, in which charginos, not neutralinos, are produced in the final state. These charginos then decay into W bosons, whose leptonic decays provide another powerful signature of Stealth SUSY. Such final states can be modeled without using the assumption of S_T -scaling, and such an analysis would complement this Thesis.




Figure 9.14: Diagrams of Stealth SUSY interactions with vectorlike portal.

Stealth SUSY with displaced jets

We mentioned in Section 2.2.2 that the assumptions of this analysis are not particularly sensitive to the portal connecting the Stealth sector with the MSSM. However, this is only true up to a point. Fig. 9.14 shows example decays with a vectorlike portal. Note, in particular, the decay of the singlet to two gluons on the right in Fig. 9.14. The lifetime of the singlet depends on the mass of the messenger that runs through the loop, and above a mass of ~ 10 TeV, the distance between the primary vertex and the location of the S decay is of order $100 \mu\text{m}$. A final state involving this decay would then have a displaced jets signature, in which one dijet in an event has an origin that is some distance from the primary vertex reconstructed for the event.

With such a final state, one would then have a three-dimensional parameter-space, with the mass of the messenger particle as the third dimension.

Bibliography

- [1] M. Srednicki. *Quantum field theory*. Cambridge University Press, Jan. 2007. ISBN: 978-0-521-86449-7, 978-0-511-26720-8 (cit. on p. 3).
- [2] Bernd A. Kniehl, Andrey F. Pikelner, and Oleg L. Veretin. “mr: a C++ library for the matching and running of the Standard Model parameters”. *Comput. Phys. Commun.* 206 (2016), pp. 84–96. DOI: [10.1016/j.cpc.2016.04.017](#). arXiv: [1601.08143 \[hep-ph\]](#) (cit. on p. 17).
- [3] Y. Ashie et al. “A Measurement of atmospheric neutrino oscillation parameters by SUPER-KAMIOKANDE I”. *Phys. Rev. D* 71 (2005), p. 112005. DOI: [10.1103/PhysRevD.71.112005](#). arXiv: [hep-ex/0501064](#) (cit. on p. 18).
- [4] Q. R. Ahmad et al. “Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory”. *Phys. Rev. Lett.* 89 (2002), p. 011301. DOI: [10.1103/PhysRevLett.89.011301](#). arXiv: [nucl-ex/0204008](#) (cit. on p. 18).
- [5] N. Aghanim et al. “Planck 2018 results. VI. Cosmological parameters”. *Astron. Astrophys.* 641 (2020). [Erratum: *Astron. Astrophys.* 652, C4 (2021)], A6. DOI: [10.1051/0004-6361/201833910](#). arXiv: [1807.06209 \[astro-ph.CO\]](#) (cit. on p. 19).
- [6] Sidney R. Coleman and J. Mandula. “All Possible Symmetries of the S Matrix”. *Phys. Rev.* 159 (1967). Ed. by A. Zichichi, pp. 1251–1256. DOI: [10.1103/PhysRev.159.1251](#) (cit. on p. 19).
- [7] Wim Beenakker et al. “NNLL-fast: predictions for coloured supersymmetric particle production at the LHC with threshold and Coulomb resummation”. *JHEP* 12 (2016), p. 133. DOI: [10.1007/JHEP12\(2016\)133](#). arXiv: [1607.07741 \[hep-ph\]](#) (cit. on pp. 32, 80).
- [8] Stephen P. Martin. “A Supersymmetry primer”. *Adv. Ser. Direct. High Energy Phys.* 18 (1998). Ed. by Gordon L. Kane, pp. 1–98. DOI: [10.1142/9789812839657_0001](#). arXiv: [hep-ph/9709356](#) (cit. on p. 35).
- [9] Gerard Jungman, Marc Kamionkowski, and Kim Griest. “Supersymmetric dark matter”. *Phys. Rept.* 267 (1996), pp. 195–373. DOI: [10.1016/0370-1573\(95\)00058-5](#). arXiv: [hep-ph/9506380](#) (cit. on p. 35).
- [10] Matteo Viel et al. “Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data”. *Phys. Rev. D* 88

- (2013), p. 043502. DOI: [10.1103/PhysRevD.88.043502](#). arXiv: [1306.2314 \[astro-ph.CO\]](#) (cit. on p. 35).
- [11] R. D. Peccei and Helen R. Quinn. “CP Conservation in the Presence of Instantons”. *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443. DOI: [10.1103/PhysRevLett.38.1440](#) (cit. on p. 36).
 - [12] “Search for the exotic decay of the Higgs boson into two light pseudoscalars with four photons in the final state in proton-proton collisions at $\sqrt{s} = 13$ TeV” (Aug. 2022). arXiv: [2208.01469 \[hep-ex\]](#) (cit. on p. 36).
 - [13] “Search for exotic Higgs boson decays $H \rightarrow \mathcal{A}\mathcal{A} \rightarrow 4\gamma$ with events containing two merged diphotons in proton-proton collisions at $\sqrt{s} = 13$ TeV” (Sept. 2022). arXiv: [2209.06197 \[hep-ex\]](#) (cit. on p. 36).
 - [14] Thomas J. LeCompte and Stephen P. Martin. “Large Hadron Collider reach for supersymmetric models with compressed mass spectra”. *Phys. Rev. D* 84 (2011), p. 015004. DOI: [10.1103/PhysRevD.84.015004](#). arXiv: [1105.4304 \[hep-ph\]](#) (cit. on p. 40).
 - [15] Thomas J. LeCompte and Stephen P. Martin. “Compressed supersymmetry after 1/fb at the Large Hadron Collider”. *Phys. Rev. D* 85 (2012), p. 035023. DOI: [10.1103/PhysRevD.85.035023](#). arXiv: [1111.6897 \[hep-ph\]](#) (cit. on p. 40).
 - [16] Albert M Sirunyan et al. “Search for supersymmetry with a compressed mass spectrum in the vector boson fusion topology with 1-lepton and 0-lepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”. *JHEP* 08 (2019), p. 150. DOI: [10.1007/JHEP08\(2019\)150](#). arXiv: [1905.13059 \[hep-ex\]](#) (cit. on p. 40).
 - [17] Glennys R. Farrar and Pierre Fayet. “Phenomenology of the Production, Decay, and Detection of New Hadronic States Associated with Supersymmetry”. *Phys. Lett. B* 76 (1978), p. 575. DOI: [10.1016/0370-2693\(78\)90858-4](#) (cit. on p. 41).
 - [18] R. Barbier et al. “R-parity violating supersymmetry”. *Phys. Rept.* 420 (2005), p. 1. DOI: [10.1016/j.physrep.2005.08.006](#). arXiv: [hep-ph/0406039](#) (cit. on p. 41).
 - [19] Peter W. Graham, Surjeet Rajendran, and Prashant Saraswat. “Supersymmetric crevices: Missing signatures of R -parity violation at the LHC”. *Phys. Rev. D* 90 (2014), p. 075005. DOI: [10.1103/PhysRevD.90.075005](#). arXiv: [1403.7197 \[hep-ph\]](#) (cit. on p. 41).
 - [20] Matthew J. Strassler and Kathryn M. Zurek. “Echoes of a hidden valley at hadron colliders”. *Phys. Lett. B* 651 (2007), p. 374. DOI: [10.1016/j.physletb.2007.06.055](#). arXiv: [hep-ph/0604261](#) (cit. on p. 41).
 - [21] Tao Han et al. “Phenomenology of hidden valleys at hadron colliders”. *JHEP* 07 (2008), p. 008. DOI: [10.1088/1126-6708/2008/07/008](#). arXiv: [0712.2041 \[hep-ph\]](#) (cit. on p. 41).
 - [22] Yuichiro Nakai, Matthew Reece, and Ryosuke Sato. “SUSY Higgs Mass and Collider Signals with a Hidden Valley”. *JHEP* 03 (2016), p. 143. DOI: [10.1007/JHEP03\(2016\)143](#). arXiv: [1511.00691 \[hep-ph\]](#) (cit. on p. 41).

- [23] Gustavo Burdman et al. “Folded supersymmetry and the LEP paradox”. *JHEP* 02 (2007), p. 009. DOI: [10.1088/1126-6708/2007/02/009](https://doi.org/10.1088/1126-6708/2007/02/009). arXiv: [hep-ph/0609152](https://arxiv.org/abs/hep-ph/0609152) (cit. on p. 41).
- [24] Peter W. Graham et al. “Displaced Supersymmetry”. *JHEP* 07 (2012), p. 149. DOI: [10.1007/JHEP07\(2012\)149](https://doi.org/10.1007/JHEP07(2012)149). arXiv: [1204.6038](https://arxiv.org/abs/1204.6038) [[hep-ph](https://arxiv.org/archive/hep)] (cit. on p. 41).
- [25] Daniele S. M. Alves, Jia Liu, and Neal Weiner. “Hiding Missing Energy in Missing Energy”. *JHEP* 04 (2015), p. 088. DOI: [10.1007/JHEP04\(2015\)088](https://doi.org/10.1007/JHEP04(2015)088). arXiv: [1312.4965](https://arxiv.org/abs/1312.4965) [[hep-ph](https://arxiv.org/archive/hep)] (cit. on p. 41).
- [26] Ulrich Ellwanger and Ana M. Teixeira. “NMSSM with a singlino LSP: possible challenges for searches for supersymmetry at the LHC”. *JHEP* 10 (2014), p. 113. DOI: [10.1007/JHEP10\(2014\)113](https://doi.org/10.1007/JHEP10(2014)113). arXiv: [1406.7221](https://arxiv.org/abs/1406.7221) [[hep-ph](https://arxiv.org/archive/hep)] (cit. on p. 41).
- [27] Ben Heidenreich and Yuichiro Nakai. “Natural Supersymmetry in Warped Space”. *JHEP* 10 (2014), p. 182. DOI: [10.1007/JHEP10\(2014\)182](https://doi.org/10.1007/JHEP10(2014)182). arXiv: [1407.5095](https://arxiv.org/abs/1407.5095) [[hep-ph](https://arxiv.org/archive/hep)] (cit. on p. 41).
- [28] JiJi Fan, Matthew Reece, and Joshua T. Ruderman. “Stealth Supersymmetry”. *JHEP* 11 (2011), p. 012. DOI: [10.1007/JHEP11\(2011\)012](https://doi.org/10.1007/JHEP11(2011)012). arXiv: [1105.5135](https://arxiv.org/abs/1105.5135) [[hep-ph](https://arxiv.org/archive/hep)] (cit. on p. 41).
- [29] JiJi Fan, Matthew Reece, and Joshua T. Ruderman. “A Stealth Supersymmetry Sampler”. *JHEP* 07 (2012), p. 196. DOI: [10.1007/JHEP07\(2012\)196](https://doi.org/10.1007/JHEP07(2012)196). arXiv: [1201.4875](https://arxiv.org/abs/1201.4875) [[hep-ph](https://arxiv.org/archive/hep)] (cit. on p. 41).
- [30] JiJi Fan et al. “Stealth Supersymmetry Simplified”. *JHEP* 07 (2016), p. 016. DOI: [10.1007/JHEP07\(2016\)016](https://doi.org/10.1007/JHEP07(2016)016). arXiv: [1512.05781](https://arxiv.org/abs/1512.05781) [[hep-ph](https://arxiv.org/archive/hep)] (cit. on p. 41).
- [31] Vardan Khachatryan et al. “Search for stealth supersymmetry in events with jets, either photons or leptons, and low missing transverse momentum in pp collisions at 8 TeV”. *Phys. Lett. B* 743 (2015), p. 503. DOI: [10.1016/j.physletb.2015.03.017](https://doi.org/10.1016/j.physletb.2015.03.017). arXiv: [1411.7255](https://arxiv.org/abs/1411.7255) [[hep-ex](https://arxiv.org/archive/hep)] (cit. on pp. 44, 96, 97, 157).
- [32] “LHC Machine”. *JINST* 3 (2008). Ed. by Lyndon Evans and Philip Bryant, S08001. DOI: [10.1088/1748-0221/3/08/S08001](https://doi.org/10.1088/1748-0221/3/08/S08001) (cit. on p. 48).
- [33] S. Chatrchyan et al. “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) (cit. on pp. 52, 61).
- [34] S Chatrchyan et al. “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). arXiv: [0910.5530](https://arxiv.org/abs/0910.5530) [[physics.ins-det](https://arxiv.org/archive/physics)] (cit. on p. 54).
- [35] Benedikt Vormwald. “The CMS inner tracker - transition from LHC Run I to Run II and first experience of Run II”. *PoS EPS-HEP2015* (2015), p. 229. DOI: [10.22323/1.234.0229](https://doi.org/10.22323/1.234.0229) (cit. on p. 55).
- [36] W. Adam et al. “The CMS Phase-1 Pixel Detector Upgrade”. *JINST* 16.02 (2021), 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) [[physics.ins-det](https://arxiv.org/archive/physics)] (cit. on p. 56).
- [37] A. Benaglia. “The CMS ECAL performance with examples”. *JINST* 9 (2014). Ed. by Pietro Govoni et al., p. C02008. DOI: [10.1088/1748-0221/9/02/C02008](https://doi.org/10.1088/1748-0221/9/02/C02008) (cit. on p. 58).

- [38] Tanmay Mudholkar. “CMS ECAL monitoring and calibration in LHC Run II”. *J. Phys. Conf. Ser.* 1162.1 (2019), p. 012002. DOI: [10.1088/1742-6596/1162/1/012002](https://doi.org/10.1088/1742-6596/1162/1/012002) (cit. on p. 60).
- [39] P. Paolucci et al. “CMS Resistive Plate Chamber overview, from the present system to the upgrade phase I”. *PoS RPC2012* (2012), p. 004. DOI: [10.1088/1748-0221/8/04/P04005](https://doi.org/10.1088/1748-0221/8/04/P04005). arXiv: [1209.1941](https://arxiv.org/abs/1209.1941) [[physics.ins-det](#)] (cit. on p. 64).
- [40] “CMS Technical Design Report for the Level-1 Trigger Upgrade” (June 2013). Ed. by A. Tapper and Darin Acosta (cit. on p. 68).
- [41] A. M. Sirunyan et al. “Particle-flow reconstruction and global event description with the CMS detector”. *JINST* 12.10 (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) [[physics.ins-det](#)] (cit. on p. 72).
- [42] Serguei Chatrchyan et al. “Description and performance of track and primary-vertex reconstruction with the CMS tracker”. *JINST* 9.10 (2014), P10009. DOI: [10.1088/1748-0221/9/10/P10009](https://doi.org/10.1088/1748-0221/9/10/P10009). arXiv: [1405.6569](https://arxiv.org/abs/1405.6569) [[physics.ins-det](#)] (cit. on p. 72).
- [43] K. Rose. “Deterministic annealing for clustering, compression, classification, regression, and related optimization problems”. *IEEE Proc.* 86.11 (1998), pp. 2210–2239. DOI: [10.1109/5.726788](https://doi.org/10.1109/5.726788) (cit. on p. 73).
- [44] P. Adzic et al. “Reconstruction of the signal amplitude of the CMS electromagnetic calorimeter”. *Eur. Phys. J. C* 46S1 (2006), pp. 23–35. DOI: [10.1140/epjcd/s2006-02-002-x](https://doi.org/10.1140/epjcd/s2006-02-002-x) (cit. on p. 74).
- [45] Serguei Chatrchyan et al. “Performance of CMS Muon Reconstruction in pp Collision Events at $\sqrt{s} = 7$ TeV”. *JINST* 7 (2012), P10002. DOI: [10.1088/1748-0221/7/10/P10002](https://doi.org/10.1088/1748-0221/7/10/P10002). arXiv: [1206.4071](https://arxiv.org/abs/1206.4071) [[physics.ins-det](#)] (cit. on p. 74).
- [46] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti- k_t jet clustering algorithm”. *JHEP* 04 (2008), p. 063. DOI: [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063). arXiv: [0802.1189](https://arxiv.org/abs/0802.1189) [[hep-ph](#)] (cit. on p. 75).
- [47] 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), p. 079. DOI: [10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079). arXiv: [1405.0301](https://arxiv.org/abs/1405.0301) [[hep-ph](#)] (cit. on p. 77).
- [48] R. Frederix et al. “The automation of next-to-leading order electroweak calculations”. *JHEP* 07 (2018). [Erratum: *JHEP* 11, 085 (2021)], p. 185. DOI: [10.1007/JHEP11\(2021\)085](https://doi.org/10.1007/JHEP11(2021)085). arXiv: [1804.10017](https://arxiv.org/abs/1804.10017) [[hep-ph](#)] (cit. on p. 77).
- [49] Richard D. Ball et al. “Parton distributions from high-precision collider data”. *Eur. Phys. J. C* 77 (2017), p. 663. DOI: [10.1140/epjc/s10052-017-5199-5](https://doi.org/10.1140/epjc/s10052-017-5199-5). arXiv: [1706.00428](https://arxiv.org/abs/1706.00428) [[hep-ph](#)] (cit. on p. 80).
- [50] Benjamin Fuks et al. “Gaugino production in proton-proton collisions at a center-of-mass energy of 8 TeV”. *JHEP* 10 (2012), p. 081. DOI: [10.1007/JHEP10\(2012\)081](https://doi.org/10.1007/JHEP10(2012)081). arXiv: [1207.2159](https://arxiv.org/abs/1207.2159) [[hep-ph](#)] (cit. on p. 80).

- [51] Benjamin Fuks et al. “Precision predictions for electroweak superpartner production at hadron colliders with Resummino”. *Eur. Phys. J. C* 73 (2013), p. 2480. DOI: [10.1140/epjc/s10052-013-2480-0](#). arXiv: [1304.0790 \[hep-ph\]](#) (cit. on p. 80).
- [52] Christoph Borschensky et al. “Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV”. *Eur. Phys. J. C* 74 (2014), p. 3174. DOI: [10.1140/epjc/s10052-014-3174-y](#). arXiv: [1407.5066 \[hep-ph\]](#) (cit. on p. 80).
- [53] Wim Beenakker et al. “NNLL resummation for squark and gluino production at the LHC”. *JHEP* 12 (2014), p. 023. DOI: [10.1007/JHEP12\(2014\)023](#). arXiv: [1404.3134 \[hep-ph\]](#) (cit. on p. 80).
- [54] S. Abdullin et al. “The fast simulation of the CMS detector at LHC”. *J. Phys. Conf. Ser.* 331 (2011). Ed. by Simon C. Lin, p. 032049. DOI: [10.1088/1742-6596/331/3/032049](#) (cit. on p. 80).
- [55] Richard D. Ball et al. “Parton distributions with QED corrections”. *Nucl. Phys. B* 877 (2013), p. 290. DOI: [10.1016/j.nuclphysb.2013.10.010](#). arXiv: [1308.0598 \[hep-ph\]](#) (cit. on p. 80).
- [56] T. Gleisberg et al. “Event generation with SHERPA 1.1”. *JHEP* 02 (2009), p. 007. DOI: [10.1088/1126-6708/2009/02/007](#). arXiv: [0811.4622 \[hep-ph\]](#) (cit. on p. 80).
- [57] Hung-Liang Lai et al. “New parton distributions for collider physics”. *Phys. Rev. D* 82 (2010), p. 074024. DOI: [10.1103/PhysRevD.82.074024](#). arXiv: [1007.2241 \[hep-ph\]](#) (cit. on p. 80).
- [58] S. Agostinelli et al. “GEANT4—a simulation toolkit”. *Nucl. Instrum. Meth. A* 506 (2003), p. 250. DOI: [10.1016/S0168-9002\(03\)01368-8](#) (cit. on p. 80).
- [59] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The Catchment Area of Jets”. *JHEP* 04 (2008), p. 005. DOI: [10.1088/1126-6708/2008/04/005](#). arXiv: [0802.1188 \[hep-ph\]](#) (cit. on p. 89).
- [60] Vardan Khachatryan et al. “Search for Microscopic Black Hole Signatures at the Large Hadron Collider”. *Phys. Lett. B* 697 (2011), p. 434. DOI: [10.1016/j.physletb.2011.02.032](#). arXiv: [1012.3375 \[hep-ex\]](#) (cit. on p. 96).
- [61] Serguei Chatrchyan et al. “Search for microscopic black holes in pp collisions at $\sqrt{s} = 7$ TeV”. *JHEP* 04 (2012), p. 061. DOI: [10.1007/JHEP04\(2012\)061](#). arXiv: [1202.6396 \[hep-ex\]](#) (cit. on p. 96).
- [62] Serguei Chatrchyan et al. “Search for Microscopic Black Holes in pp Collisions at $\sqrt{s} = 8$ TeV”. *JHEP* 07 (2013), p. 178. DOI: [10.1007/JHEP07\(2013\)178](#). arXiv: [1303.5338 \[hep-ex\]](#) (cit. on p. 96).
- [63] Albert M Sirunyan et al. “Search for black holes in high-multiplicity final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”. *Phys. Lett. B* 774 (2017), p. 279. DOI: [10.1016/j.physletb.2017.09.053](#). arXiv: [1705.01403 \[hep-ex\]](#) (cit. on p. 96).

- [64] Kyle S. Cranmer. “Kernel estimation in high-energy physics”. *Comput. Phys. Commun.* 136 (2001), p. 198. DOI: [10.1016/S0010-4655\(00\)00243-5](https://doi.org/10.1016/S0010-4655(00)00243-5). arXiv: [hep-ex/0011057](https://arxiv.org/abs/hep-ex/0011057) (cit. on p. 98).
- [65] Jerzy Neyman and Egon Sharpe Pearson. “On the Problem of the Most Efficient Tests of Statistical Hypotheses”. *Phil. Trans. Roy. Soc. Lond. A* 231.694-706 (1933), pp. 289–337. DOI: [10.1098/rsta.1933.0009](https://doi.org/10.1098/rsta.1933.0009) (cit. on p. 136).
- [66] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. *Annals Math. Statist.* 9.1 (1938), pp. 60–62. DOI: [10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360) (cit. on p. 137).
- [67] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. *Eur. Phys. J. C* 71 (2011). [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)], p. 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) [[physics.data-an](https://arxiv.org/archive/data)] (cit. on pp. 137, 140).
- [68] Thomas Junk. “Confidence level computation for combining searches with small statistics”. *Nucl. Instrum. Meth. A* 434 (1999), p. 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) (cit. on p. 140).
- [69] Alexander L. Read. “Presentation of search results: The CL(s) technique”. *J. Phys. G* 28 (2002). Ed. by M. R. Whalley and L. Lyons, p. 2693. DOI: [10.1088/0954-3899/28/10/313](https://doi.org/10.1088/0954-3899/28/10/313) (cit. on p. 140).
- [70] The ATLAS Collaboration, The CMS Collaboration, The LHC Higgs Combination Group. *Procedure for the LHC Higgs boson search combination in Summer 2011*. CMS Note CMS-NOTE-2011-005. CERN, 2011. URL: <https://cds.cern.ch/record/1379837> (cit. on p. 140).