

A SEARCH FOR RESONANT CASCADES TO NEUTRAL STANDARD MODEL
BOSONS AND P_T^{MISS} AND SELECTED STORIES OF THE CMS HADRON
CALORIMETER PHASE 1 UPGRADE

Grace Elaine Cummings
Selbyville, DE

Bachelor of Science, Virginia Commonwealth University, 2016

A Dissertation submitted to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of
Doctor of Philosophy

Department of Physics

University of Virginia
December 2022

Robert J. Hirosky, Chair

Robert C. Group

Diana Vaman

Bradley Johnson



A search for resonant cascades to neutral Standard Model bosons and p_T^{miss} and selected stories of the CMS Hadron Calorimeter Phase 1 upgrade

Grace Elaine Cummings

(ABSTRACT)

Direct decays of proposed heavy force mediator particles to standard model (SM) leptons have been excluded to high masses [1], but more complex scenarios remain unexplored. This thesis presents a search for a leptophobic Z' decaying via anomalous in proton-proton collisions at $\sqrt{s} = 13$ TeV with the Compact Muon Solenoid (CMS) Experiment. A leptophobic Z' can decay via a pair of anomalous, new beyond the standard model (BSM) fermions introduced to cancel the gauge anomalies arising from the leptophobic condition. These heavy intermediate particles decay in turn to neutral SM bosons and lighter, stable anomalous. The stable anomalous can serve as a dark matter candidate. This analysis targets the $ZH \rightarrow \mu\bar{\mu}b\bar{b} + \cancel{p}_T^{miss}$ final state in a total integrated luminosity of 137.6 fb^{-1} corresponding to the 2016-2018 data set. To search for the resonant Z' production, this analysis employs Recursive Jigsaw Reconstruction (RJR), an iterative framework to reconstruct mass estimators in systems with invisible particles in the final state. This thesis presents the expected sensitivity of this novel model and observable in a CMS search. Additionally, the assembly, installation, and commissioning of CMS Hadron Calorimeter Phase 1 upgrade is discussed.

Dedication

To Sussex County, Delaware.

Acknowledgments

My PhD has been my entire education in the making. I have to thank my high school teachers Ms. Marjorie Kimmel and Dr. Corey Heacock. When the LHC started, Ms. Kimmel made sure my ninth grade class knew what it was and knew it was cool. Dr. Heacock gave me the first taste of academic rigor. Not only did her class prepare me for college, she was the first person to introduce me to the CMS Experiment.

My community in undergrad gave me the tools and perspective to succeed. The inspiration of my friends and the TEDxVCU and Indie Lab teams made me the capable person I am today. Academically, Dr. McMullen and Dr. Reiner's courses and offered research experiences established my work ethic and prepared me for what was to come.

The help and support I have received at UVa cannot be fully captured here, but I can try. I need to thank my friends and the rest of my cohort (and some of those from the years above) who kept me afloat. Charles Glazer's help was indispensable, and deserves special recognition. Thomas Anderson's technical prowess paved my way to some modicum of hardware savvy. Diana Vaman and Simonetta Liuti were my tireless advocates as my research committee and cheered me on through the pandemic darkness. Thank you to Diana, Craig Group and Brad Johnson, for serving on my thesis committee and reading this document. When I eventually returned to UVa, Thomas Krahulik gave this CERN vagabond a place to stay while she finished writing and he weathered that with grace and unflinching compassion.

Of course, I owe the largest thanks to my adviser Bob. He always made time for me, was not afraid to get in the thick of it, and supported me through all of my too many endeavors. He rarely told me no, for better or worse. Thank you for truly helping me become the best I can be.

Professionally, I need to thank the analysis team for their continual support: John, Tongguang, Bob, and Atanu. And a special thanks to John for always entertaining my questions, and incessantly keeping me on my toes with his own. All of you made this physics analysis possible, and for that I thank you. Furthermore, the importance of the casual conversations and the official CMS review cannot be understated. It truly takes a village.

The HCAL team was my first true CMS community. Starting with my work at Fermilab, I was launched into the roles that would make my name. Through the

ngCCM team (first round), I had the pleasure of working – and ultimately becoming friends – with Furkan, Fatma, Markus, and Francesca. Not only were they an amazing team, they valiantly supported me through my surgery and recovery, and for that I am eternally grateful. During HCAL installation, Alexey Kalinin and Jay Lawhorn were rock stars that taught me everything I needed to know. Through the VTRx investigation, I had the pleasure of working closely with Jay Dittmann and Sebastian. I could not have been luckier to have been paired during the pandemic with folks that were both excellent scientists and kind souls. This is just a small subset of the HCAL team, all of which I am indebted too, and grateful for.

I spent nearly my entire PhD research career based at CERN which is a privilege I do not take lightly. The folks I have met through my time in Geneva, Switzerland have not only taught me innumerable things about physics, but have become some of my best friends. Many of us were together during the COVID-19 lockdowns, so I want to give a special thank you to the “Quaranteam”: Aleksandra, Sam, Sophie, and Kwame. We were friends before, but we were family after. My roommates were there through it all, and I think we solved many of the world’s (and physics’) problems at our kitchen table; thank you for making our house a home. Finally, a thank you to Tal. You made the best of the worst time and still somehow managed to talk me off the ceiling long enough to do some physics. This does not touch at all upon all of the colleagues and friends I passed in the hall, shared offices with, took shifts with, or were the friends of friends or occasional visitors that made life rich. CERN became a second home, and for that I am forever grateful.

Geneva and the Pays de Gex became my home in no small part due to the love and hospitality of Pascal Bulteau and Pascale Brunet. You both welcomed me into your home, shared your culture, and taught me your language. I am sorry I did not add the appendix of French cheeses, but I carry that knowledge with me now. Your support was unending, and I am forever grateful for it. Thank you for taking me in during Christmas 2020. You were there for me when I needed it most.

Finally, I need to thank my family. I need to thank my Mom-mom, who did not get to see me finish, but from whom I learned academic discipline and curiosity. I have enough degrees for both of us now. Daisy anchored me to home, and continually reminded me I was loved. And I have to thank my parents: Mom, Dad, Craig, and Cyndi. Your pride, support, and seemingly endless trips to airports/universities during rush hour traffic have finally paid off.

Contents

List of Figures	viii
List of Tables	ix
1 Introduction	1
2 Theoretical Underpinnings	3
2.1 The Standard Model of particle physics	3
2.1.1 Particles	3
2.1.2 Interactions	5
2.1.3 Limitations and Open Questions	18
2.2 Heavy vector bosons: A Common Beyond the Standard Model Ingredient	21
2.2.1 Overview of Heavy Vector Bosons	22
2.2.2 Leptophobic cascade decay	24
3 The CMS Experiment and the LHC	29
3.1 The Large Hadron Collider	29
3.1.1 Beam Properties	30
3.1.2 Accelerator Chain and LHC Details	31
3.1.3 Delivered Beam Characteristics	34
3.2 The CMS Experiment	35
3.2.1 Coordinate System	37
3.2.2 Tracker	38
3.2.3 ECAL	40
3.2.4 HCAL	41
3.2.5 Magnet and Flux Return Yolk	43

3.2.6	Muon System	44
3.2.7	Trigger System	48
4	The Hadron Calorimeter Phase 1 Upgrade	50
4.1	HCAL Endcap and Barrel Overview, and Upgrade Motivation	51
4.1.1	HBHE Segmentation and Nomenclature	51
4.1.2	Phase 0 HCAL Challenges	55
4.1.3	HL-LHC Demands	55
4.1.4	Phase 1 Upgrade Overview	55
4.2	Lock, Stock, and HCAL Barrel	57
4.2.1	Phase 1 Barrel Front End	57
4.2.2	Next Generation Clock Control and Monitoring	58
4.2.3	HCAL Barrel Installation	63
4.3	VTRx Investigations	66
4.3.1	Description and extent of the phenomenon	66
4.3.2	VTRx Problem: HCAL Solution	69
4.3.3	Cooling Fin Properties	70
4.3.4	Problem origin	71
4.3.5	Final Remarks	73
5	Particle Reconstruction	74
5.1	The Particle Flow Algorithm	74
5.2	Muons	78
5.3	Electrons	78
5.4	Jets	81
5.5	Missing Transverse Momentum	83
6	Strategy and Event Selections	85
6.1	Model Phenomenology and Final State	85
6.2	Analysis Strategy	86

6.3	Recursive Jigsaw Reconstruction	91
6.4	Object Specifics	93
6.4.1	Electrons	96
6.4.2	Muons	97
6.4.3	Dimuon Z Candidate	97
6.4.4	Jets	98
6.4.5	Higgs Candidate Reconstruction	103
6.4.6	Missing Transverse Momentum	103
6.5	Preselection Cut Scheme	103
7	Background Estimations	109
7.1	$t\bar{t}$ Estimation	109
7.1.1	$e\mu$ channel normalization	111
7.1.2	$e\mu$ -channel informed estimation of $\mu\mu$ background	115
7.2	α -Method for $Z + \text{jets}$ Estimation	116
7.2.1	Background Normalization	117
7.2.2	RJR Background Shape Parameterization and Estimation	118
7.2.3	α -method validation	125
7.3	Diboson Estimation	125
8	Systematic Uncertainties	127
8.1	$Z + \text{jets}$ Estimation with α -Method	127
8.2	Data-driven $t\bar{t}$ estimation	129
8.3	Jet Energy Scale	134
8.4	b-tagging	134
8.5	p_T^{miss} Uncertainties	137
8.6	Muon ID	141
8.7	Muon Trigger Scale Factor	141
8.8	HEM15/16	142

8.9	Parton Distribution Function (PDF) Scale	145
8.10	QCD renormalization and Factorization scale	145
9	Statistical Interpretation and Results	147
9.1	Statistical Analysis	147
9.2	Input Model	149
9.2.1	Signal Interpolation	152
9.3	Checks of Statistical Model	152
9.3.1	Pulls	153
9.3.2	Impacts	155
9.3.3	Signal Injection Tests	155
9.4	Results	164
10	Conclusions and Outlook	168
	Bibliography	171
	Appendices	180
	Appendix A Scale Factors	181
A.1	Muon Scale Factors	181
A.2	Electron Scale Factors	185
A.3	Btagging Scale Factors	185
	Appendix B Control Plots for the electron and muon data set used to estimate $t\bar{t}$bar - Preselection and Jet Mass Signal Region	189
	Appendix C α-method validation region closure plots	195
	Appendix D Impact of low background on Signal Injection Tests	196

List of Figures

2.1	Summary of the SM particles and their properties. [5, 6, 7]	4
2.2	Potential for spontaneously breaking a $N = 2$ symmetry. Minima would correspond to the vacuum expectation value. [8]	14
2.3	Summary of current public SM EW interaction cross section measurements by the Compact Muon Solenoid Experiment. [15]	19
2.4	Summary of the Compact Muon Solenoid Experiment’s “Beyond Two (2) Generations” (B2G) physics analysis group 2016, 2017, and 2018 Z' results [34]. The B2G group is a BSM search group in CMS that specializes in intermediate vector bosons in the final state.	24
2.5	Summary of the Compact Muon Solenoid Experiment’s “Exotica” (EXO) physics analysis group 2016, 2017, and 2018 Z' results [35]. The EXO group is a BSM search group in CMS that searches generally for exotic signatures.	24
2.6	Triangle diagrams of an axial current vertex.	25
2.7	Illustrative decays of the baryonic Z' to anomalous.	28
3.1	The LHC accelerator chain and the experiments.[42]	32
3.2	Picture of cryodipole outside of the CERN Restaurant 1 patio.	33
3.3	Total delivered and recorded integrated luminosity (3.3a) and average number of pile-up interactions (3.3b) measured at CMS during Run 2 of the LHC. Blank periods in the integrated luminosity plot show the periods of Year End Technical Stop (YETS) when the machine is down for maintenance.[47]	35
3.4	A diagram of the relationship between η and θ . Pseudorapidity values mirrored across the z -axis are negative.[50]	38
3.5	A cross-section schematic of the Phase 0 CMS tracker.[44]	40
3.6	A cross-section comparison of the Phase 0 CMS pixel tracker (turquoise,bottom) and Phase 1 pixel tracker (blue, top) .[51]	40
3.7	Labeled locations of the the HCAL subdetectors.	42

3.8	A photo of a yoke endcap disk, the central wheel with the vacuum tank, and an additional barrel wheel during magnet assembly in SX5.[44]	44
3.9	Left: Stylized diagram of a functioning cross-section of a drift tube. Right: Stylized diagram of a CSC showing the azimuthal wires, and the perpendicular strips.[44]	46
3.10	Quarter-view of CMS along $z - y$ to show the location of the CSCs. The CSCs are highlighted in red, and are label “ME”X.[44]	47
3.11	Example CMS barrel wheel in $x - y$ showing the locations of the RPCs and DTs [44].	47
3.12	Diagram of flow of L1 trigger.[44].	49
4.1	The numbering begins on the inside edge of the LHC ring. For example, looking at HBM while standing on a cherry picker in front of HEM, HBM1 would be at 9 o’clock, and HBM10 would be a 3 o’clock. [44]	52
4.2	A diagram of an HCAL Barrel wedge to show staggered layers [44].	52
4.3	An yz -plane slice of the Phase 0 HCAL Barrel and Endcap. Towers and scintillating layers are indicated. Colored regions show the summed readout per-tower in the Phase 0 system. “FEE” indicates the location of the front end electronics, or the “RBx”.	53
4.4	An image of HBP prior to insertion into the magnet’s vacuum tank above ground in SX5 in March, 2006. [54]. Image shows not only the scale of one half of the HCAL Barrel, but the layered structure of the absorber. This is no longer visible, as it has been covered by the cable trays of the tracker.	54
4.5	An yz -plane slice of the Phase 1 HCAL Barrel and Endcap. Towers and scintillating layers are indicated. Colored regions show the summed readout per-tower of the upgraded Phase 1 system. “FEE” indicates the location of the front end electronics, or the “RBx”.	56
4.6	A schematic showing the naming, constituents, and organization of an HB RBx. Since HB RBxs are in 360°C , this specific orientation would only be found in the 6 o’clock position of either HB half.	58
4.7	A block diagram of an HB ngCCM. Only one control card interfaces with the backend concurrently.	59

- 4.8 All components of a ngCCM, partially disassembled. Module is constructed left to right. The clock card is already installed between the bottom and a middle layer of housing in the leftmost part of the image. Additional copper heat sinks, called cooling fins, are also installed over the VTRx. The blue thermal gap pads used are visible on the heat sinks. 60
- 4.9 (a) Cooling fin installation on the VTRx from a view of the face of the control card. The module is partially disassembled to reveal fin. The cooling fin is in physical contact with the optical sub-assemblies of the VTRx via a 7 mm x 3 mm 0.1 in thick Bergquist thermal gap pad. The VTRx transmitter ASIC is also in contact with the larger plane of the cooling fin with a 12 mm x 12 mm 0.06 in thick Bergquist thermal gap pad (not visible). (b) Temperature measurements for the VTRx in an HB ngCCM. Measurements were taken with Capteur RTD PT100 PR PRO probes, and aluminum housing measurements were verified against internal ngCCM temperature sensors. 61
- 4.10 Photo of the ngCCM teststand in use with a fully assembled HB ngCCM with the face plate installed. The purple PCB is the front end emulator. The green PCB is the JTAG test board, currently in use. The debug interface is not connected. The blue optical fiber protruding from the ngCCM is connected to the primary control card, and the edge of the Glibv2 can be seen to the left side, on top of the power supply. Additional HB ngCCM can be seen. 62
- 4.11 The left image shows the minus-side Surkov frame, connected via the passerelle. The right image was taken from the central platform inside the frame looking out of the frame, with the endcap nose in view. . . 64
- 4.12 In (a) Alexey Kalinin accesses a bottom sector. The image in (b) shows the author braced to access a side-sector in the barrel. The image was taken from the floor of the vacuum tank. The sector accessed in (c) can only be accessed from a crouch from a beam, and the top sectors (d) are hard due to fighting gravity, and because one cannot stand-up straight. 65
- 4.13 Image (a) shows two connected RMs through the access hole. The thick black fibers are the megatile fibers. The black tubes with metal connectors are used for calibration. The light blue fibers are control fibers. The bright blue tubes and the red cable are the dry gas and high voltage, respectively. Image (b) shows a close-up of a partially connected RM. The Layer 0 connector (right, without fiber) has an opaque reflection, showing a damaged RM optic. Image (c) shows the organized megatile fibers from the megatile-connection end. 67

4.14	RSSI plot for all active links in the affected endcap over the lost module's corresponding drift period. The HEM09 and HEP06 deviations are clear in comparison to the relative stability in the other modules.	68
4.15	Image of fiber face after removal from drifting VTRx (left). Outgassed material appears as a white substance causing the perceived texture on the fiber face, and obscuring the central part of the fiber where light is transmitted(see red circle). Image of clean fiber for comparison (right).Red circled region is clear.	72
4.16	Mapping of operational phase space taken from measurements of a single, illustrative VTRx experiencing RSSI drift. This shows how the temperature difference between the fiber ferrule and the VTRx receiver influences the condensation of the outgassed material. The z -axis is the derivative of the RSSI, with a negative trend indicating the RSSI was decreasing, and a positive trend indicating that the RSSI was increasing. The trend in the RSSI reveals if the temperature settings either encourage condensation, or are unfavorable to condensation. Measurements are not available for the gray region.	73
5.1	Transverse section of CMS illustrating locations of different particle interactions. [63]	75
6.1	Cartoon to illustrate two general classes of the model's decay topologies. The left figure illustrates the scenario of nearly-at-rest production of heavy N_D particles. The right figure illustrates boosted N_D production, resulting in back-to-back signatures. This analysis prioritizes the former configuration where the SM backgrounds are strongly suppressed.	87
6.2	Theoretical cross sections for the $m_{N_S} = 200$ GeV signal scenario with $g_{Z'} = 0.4$ and assuming a 100% branching fraction. The cross section peaks between $m_{Z'}$ values of 3500 GeV and 4000 GeV.	88
6.3	Grid of mass points generated for signal production.	89
6.4	Overview of recursive jigsaw reconstruction. The left image shows the decay tree and rest frames for the reconstruction analysis. Example of correlation between mass estimator and input $m_{Z'}$ for various working points (arbitrary normalization).	93
6.5	Accuracy of RJR Z' mass estimator ($M_{Z'}$) using the Recursive Jigsaw Reconstruction scenario JR V.1, JR VI.1, and JR VI.3, the Contra-Boost Invariant. The greener the better. The top plot is for a signal sample grid with a fixed m_{N_S} of 200 GeV. The bottom is a signal sample grid with a fixed m_{N_S} of 1 GeV.	94

- 6.6 Accuracy of RJR Z' mass estimator ($M_{Z'}$) using the Recursive Jigsaw Reconstruction scenario JR V.1, JR VI.1, and JR VII.3, the Invisible Minimize ΔM . The greener the better. The top plot represents a signal sample grid with a fixed m_{N_S} of 200 GeV. The bottom plots is for a signal sample grid with a fixed m_{N_S} of 1 GeV. In comparison to Fig. 6.5, the only region with improved accuracy is the $m_{N_S} = 1$ GeV region with an $(m_{N_D} - m_{N_S}) < 500$ GeV. 95
- 6.7 The figure on the left shows the generator-level ΔR distribution at preselection for the Z and H in select signal samples. The middle and right figures show the ΔR distribution at preselection of the reco level leading (middle) or subleading (right) muon and the generator-level H. The overlap region corresponds with $\Delta R < 0.8$. Plots with reconstruction level objects were taken from samples with reclustering performed. The preselections applied are summarized in Section 6.5. 99
- 6.8 Kinematic plots of the $m_{Z'}$ 3000 GeV m_{N_D} 1200 GeV m_{N_S} 1 GeV signal mass point at skim level selections, in the Z candidate and H candidate overlap region ($\Delta R(Z, H) < 0.8$). Skims selections require the Z candidate reconstruction, $p_T^Z > 150.0$ GeV, and at least one fat jet. H candidate chosen as the fat jet closest to the Higg's mass. Distributions are normalized to unit area to show shape differences. 100
- 6.9 ΔR distribution between the Z candidate and H candidate at preselection in the H channel signal region for a few benchmark signal mass distributions. Scaled to full Run 2 luminosity, and a cross section of 10 fb^{-1} . Left figure is reclustered. Right figure is not reclustered. 101
- 6.10 Comparison of jet tagger efficiency for double-b or Higgs taggers. Background efficiency is defined as the weighted sum of background events passing all selections and b-tagging over the weighted sum of background events passing selections alone. Signal efficiency is defined the same way. The signal sample used is a privately generated $m_{Z'} = 2000$ GeV, $m_{N_D} = 500$ GeV, $m_{N_S} = 200$ GeV sample. The selections used were $p_T(Z) > 200$ GeV, $p_T(H) > 300$ GeV, and $p_T^{\text{miss}} > 200$ GeV. 102
- 6.11 The cutflow of total background in the sideband and the signal region for different Z candidate p_T and p_T^{miss} cut values. The H candidate has a $p_T > 300$ GeV. The grid was chosen to maintain the Z + jets and $t\bar{t}$ yields necessary to maintain the Z + jets estimation method. Roughly 200 background events are needed in the signal region. 105

6.12	Percent difference between expected limits for the cut scenarios Z candidate $p_T > 100$ GeV, $p_T^{\text{miss}} > 75$ GeV and Z candidate $p_T > 100$ GeV, $p_T^{\text{miss}} > 50$ GeV. Negative values indicate that the limit of the $p_T^{\text{miss}} > 75$ GeV scenario is lower (more sensitive) than the limit of the $p_T^{\text{miss}} > 50$ GeV.	105
6.13	The soft drop mass sideband region plots of select final state particle object kinematic distributions. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.	107
6.14	The soft drop mass sideband plots at preselection featuring the RJR mass estimator spectra. The plots are in units of GeV. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.	108
6.15	The soft drop mass signal region particle object p_T spectra. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.	108
6.16	The soft drop mass signal region plots at preselection featuring the RJR mass estimator spectra. Preselections are described in Section 6.5. The plots are in units of GeV. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.	108
7.1	Illustrative Feynman diagrams of leptonic (left) and semileptonic (right) $t\bar{t}$ decay channels. In the leptonic-case, the two leptons from the leptonic W decay fake the Z candidate and the two b-jets merge to fake the Higgs candidate jet. In the semileptonic case, a lepton from the leptonic W fakes one of the leptons in the Z candidate, but the second lepton comes from the decay of the b-quark via a virtual W boson. The hadronic top produces the fat jet candidate.	110
7.2	Value of $e\mu/\mu\mu$ in simulation versus cut for leptonic and semileptonic $t\bar{t}$. Overall, the $e\mu$ and $\mu\mu$ channels of $t\bar{t}$ respond similarly to the selections, producing a flat cutflow. This changes with the DeepAK8 tagger requirement in the leptonic channel. This produces about a 20% difference in efficiency between the $e\mu$ and $\mu\mu$ channels.	112
7.3	Distribution of ΔR between the leading lepton in the $e\mu$ fake “Z” candidate and the selected Higgs boson candidate fat jet. Semileptonic $t\bar{t}$ is back-to-back, while leptonic $t\bar{t}$ has events where the dilepton and the fat jet overlap.	113

7.4	A comparison of the shapes of the $t\bar{t}$ MC in the $e\mu$ and $\mu\mu$ channels, with the normalization scale between them. This yield is expected based on the relative inefficiency between muon and electron identification, and the make-up of our control region.	115
7.5	Overlays of the normalized $e\mu$ extrapolation (data set with scaling) and the $\mu\mu$ channel MC for preselection cuts. The plot on the left is the soft drop mass distribution. The right plot is the $t\bar{t}$ background estimation for the recursive jigsaw mass estimator of the Z' in the soft drop mass signal region.	116
7.6	Illustration of regions used in the α -method. All considered regions have a fat jet with $p_T > 300.0$ GeV and a DeepAK8 working point > 0.8 , the Z candidate $p_T > 100.0$ GeV, and an $ \eta < 2.4$. The SR label indicates the signal region, SB the sideband, VR the validation region, and ZR the blinded region around Z peak.	118
7.7	Individual backgrounds in the soft drop mass spectrum, and their respective fits. Distributions shown at preselection with the p_T^{miss} cut released, with no SR and SB separation.	119
7.8	Stacked plots showing the full soft drop mass sideband fit. Left stack has the $Z + \text{jets}$ plotted without the derived normalization, the right stack has the normalization applied. Both plots include the full background template fit to MC only, and the data sideband fit.	120
7.9	Distributions used in α -method background estimation. The lower range of the fitted region implies a final cut at 1400. The red bands are the fit uncertainties. Top Left: Sideband Drell-Yan + jets MC RJR Z' mass estimator distribution. Top Middle: RJR mass estimator sideband of $t\bar{t}$. Top Right: WZ and ZZ sideband RJR mass estimator. Bottom Left: Signal region Drell-Yan + jets MC RJR mass estimator distribution. Bottom Right: α ratio formed from the ratio of the fits in the bottom left and top left plots. Bottom Middle: Data sideband and fit.	121
7.10	Plots showing the fit agreement with the data sideband. The left plot shows the sideband data overlaid with the individual background fits. The right plot shows the agreement between the sideband data and the fit.	122
7.11	The subtracted data sideband distribution. The black line represents the $Z + \text{jets}$ contribution in the data sideband. This plot is the difference between the data sideband fit, and the $t\bar{t}$ and VV fits. Uncertainty bands are discussed in Chapter 8.	123

- 7.12 The signal region $Z + \text{jets}$ estimation. Left plot shows the estimation on a linear scale in the RJR mass estimator region between 1400 and 3000 GeV. Sideband data events exist in this region. The left plot shows the entire extrapolation range. Drell-Yan MC is included to guide the eye, and is not expected to agree. At high p_T , we rely on data, rather than simulation, to inform the shape. Uncertainty bands are discussed in Chapter 8. 124
- 7.13 Top Left: Sideband Drell-Yan + jets MC RJR Z' mass estimator distribution with derived normalization. Top Middle: Sideband of $t\bar{t}$. Top Right: WZ and ZZ sideband. Bottom Left: validation region $Z + \text{jets}$ MC. Bottom Right: α ratio formed from the ratio of the fits in the bottom left and top left plots. Bottom Middle: data sideband and fit. 125
- 7.14 Plots showing the application of the α -method in the validation region. The left plot shows the fit to the sideband data with the $t\bar{t}$ and VV components subtracted. The middle plot is the α ratio between the sidebands and the validation region. The right plot is the estimation of the $Z + \text{jets}$ contribution in the validation region through the multiplication of the functions in the two plots to the left. MC included to guide the eye. Perfect agreement between data and MC is not expected in this data-driven background. 126
- 8.1 Left: Drell-Yan signal region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted **up** in the decorrelated space. Right: Same distributions but for the Drell-Yan sideband. These functions are used to build the α ratio. The uncertainties are small, but are noticeable when propagated through the method. 130
- 8.2 Left: Drell-Yan signal region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted **down** in the decorrelated space. Right: Same distributions but for the Drell-Yan sideband. These functions are used to build the α ratio. The uncertainties are small, but are noticeable when propagated through the method. 130
- 8.3 Left: data sideband region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted **up** in the decorrelated space. Middle: Same description but for the $t\bar{t}$ MC contribution. Right: Same description but for the diboson contribution. The $t\bar{t}$ and diboson functions are subtracted from the data sideband distribution to form the distribution that gets multiplied by the α ratio. 131

8.4	Left: data sideband region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted down in the decorrelated space. Middle: Same description but for the $t\bar{t}$ MC contribution. Right: Same description but for the diboson contribution. The $t\bar{t}$ and diboson functions are subtracted from the data sideband distribution to form the distribution that gets multiplied by the α ratio.	131
8.5	Left: Different α ratio that are generated by varying one of the input function's parameters up . Middle: Different subtracted data sideband functions generated by varying one of the input function's parameters up . Right: The extrapolation to the signal region when one of the input parameters has been varied up . See Figs. 8.7 and 8.8 for larger range, logarithmic scale versions of the middle and right plots.	132
8.6	Left: Different α ratio that are generated by varying one of the input function's parameters down . Middle: Different subtracted data sideband functions generated by varying one of the input function's parameters down . Right: The extrapolation to the signal region when one of the input parameters has been varied down . See Figs. 8.7 and 8.8 for larger range, logarithmic scale versions of the middle and right plots.	132
8.7	Subtracted sideband distributions. Same as the middle plots in Figs. 8.5 and 8.6, but with extended range and logarithmic scale.	133
8.8	Shifted extrapolations with log scale and extended range. Same as the right-most plots in Figs. 8.5 and 8.6, but with extended range and logarithmic scale.	133
8.9	Plots showing the application of the up/down systematic uncertainties for the JES. The left plot is Z + jets with the full α -method applied, and the plot is the diboson MC.	135
8.10	Top: JES up/down shifts for the signal mass point $m_{Z'} 2000 m_{N_S} 400 m_{N_S} 200$. Middle: JES up/down shifts for signal mass point $m_{Z'} 4000 m_{N_D} 800 m_{N_S} 200$. Bottom: JES up/down shifts for signal mass point $m_{Z'} 5500 m_{N_D} 1800 m_{N_S} 200$. All plots are shown in the signal region.	136
8.11	Plots showing the application of the up/down systematic uncertainties for b-tagging. The left plot is Z + jets with the full α -method applied, and the plot is the diboson MC.	137
8.12	Top: btag up/down shifts for the signal mass point $m_{Z'} 2000 m_{N_S} 400 m_{N_S} 200$. Middle: btag up/down shifts for signal mass point $m_{Z'} 4000 m_{N_D} 800 m_{N_S} 200$. Bottom: btag up/down shifts for signal mass point $m_{Z'} 5500 m_{N_D} 1800 m_{N_S} 200$. All plots are shown in the signal region.	138

8.13	Plots showing the application of the up/down systematic uncertainties for unclustered MET. The left plot is Z + jets with the full α -method applied, and the plot is the diboson MC.	139
8.14	Top: Unclustered MET up/down shifts for the signal mass point $m_{Z'} 2000 m_{N_S} 400 m_{N_S} 200$. Middle: Unclustered MET up/down shifts for signal mass point $m_{Z'} 4000 m_{N_D} 800 m_{N_S} 200$. Bottom: Unclustered MET up/down shifts for signal mass point $m_{Z'} 5500 m_{N_D} 1800 m_{N_S} 200$. All plots are shown in the signal region.	140
8.15	2018 MC input distributions to the alpha method for the Z + jets estimation. Entire 2018 luminosity scaled, so any deviations over-estimate their expected value. Deviations are well within the fit uncertainties for the method, especially once combined with 2016 and 2017, as is done.	143
8.16	2018 MC input distributions for the diboson background. Entire 2018 luminosity scaled, so any deviations over-estimate their expected value. Deviations are negligible, especially once combined with 2016 and 2017, as is done.	143
8.17	Observable signal distributions in the dimuon channel with the HEM scaling applied. Entire sample scaled, so any deviations over-estimate their expected value. Deviations are negligible, especially once combined with 2016 and 2017, as is done.	144
9.1	Relationship between the RJR Z' mass estimator and select kinematic inputs. Relationship is shown for the $m_{Z'} 5500, m_{N_D} 1800$, and $m_{N_S} 200$ signal mass point in the signal region. No cross section scaling is applied.	150
9.2	The $t\bar{t}$ background estimation, taken from the electron and muon final-state $t\bar{t}$ data.	151
9.3	The diboson background estimation, taken from Monte Carlo.	151
9.4	Highest $m_{Z'}$ and m_{N_D} signal mass point with the final statistical model binning.	151
9.5	Cartoon of the relationship between the q_{ith} quantiles and the plane in $mass_1 mass_2$ -plane used to interpolate the quantile position for new mass points. In the figure, $mass_1$ and $mass_2$ are arbitrary stand-ins for $m_{Z'}$ and m_{N_S} in our signal model.	153

9.6	Selected values of fixed m_{N_D} and $m_{Z'}$ comparing the Monte Carlo generated signal samples and the interpolated samples (bold lines). The left column is at fixed m_{N_D} , and the right column is at fixed $m_{Z'}$. All plots are for the $m_{N_S} = 200$ GeV case. The interpolation appears to be correctly modeling the shape and normalization trends.	154
9.7	Signal region distributions for the example signal points used to test the statistical model.	155
9.8	Nuisance parameter pulls for fits to 500 toys with an expected signal strength equal to the median expected limit for mass point $m_{Z'}$ 2000 m_{N_D} 400 m_{N_S} 200 GeV. Pre-Fit line indicates the bare nuisance parameter value, not a pull.	156
9.9	Nuisance parameter pulls for fits to 500 toys with an expected signal strength of 0 for mass point $m_{Z'}$ 4000 m_{N_D} 800 m_{N_S} 200 GeV. Pre-Fit line indicates the bare nuisance parameter value, not a pull.	156
9.10	Nuisance parameter pulls for fits to 500 toys with an expected signal strength of 0 for mass point $m_{Z'}$ 5500 m_{N_D} 1800 m_{N_S} 200 GeV. Pre-Fit line indicates the bare nuisance parameter value, not a pull.	157
9.11	Nuisance parameter pulls for fits to 500 toys with an expected signal strength equal to the median expected limit for mass point $m_{Z'}$ 5500 m_{N_D} 1800 m_{N_S} 200 GeV. Pre-Fit line indicates the bare nuisance parameter value, not a pull.	157
9.12	Impact of nuisance parameters for the background only hypothesis with a signal hypothesis of $m_{Z'}$ 2000 GeV, m_{N_D} 400 GeV, m_{N_S} 200 GeV. The Δr is the change in the extracted signal strength with respect to the nominal when the nuisance parameter is set to its $\pm 1\sigma$ value and the fit repeated.	158
9.13	Impact of nuisance parameters for the background only hypothesis with a signal hypothesis of $m_{Z'}$ 5500 GeV, m_{N_D} 1800 GeV, m_{N_S} 200 GeV. The Δr is the change in the extracted signal strength with respect to the nominal when the nuisance parameter is set to its $\pm 1\sigma$ value and the fit repeated.	159
9.14	Trend of signal strength (r) measured versus the signal strength injected. Left: signal sample with $m_{Z'} = 2000, m_{N_D} = 400, m_{N_S} = 200$. Middle: signal sample with $m_{Z'} = 4000, m_{N_D} = 800, m_{N_S} = 200$. Right: signal sample $m_{Z'} = 5500, m_{N_D} = 1800, m_{N_S} = 200$	160

9.15	Differences (left) and pulls (right) for the measured and expected signal strengths for $m_{Z'} = 2000$, $m_{N_D} = 400$, $m_{N_S} = 200$. First row has an injected signal of 0, the second has the lower expected limit, the third the median expected limit, and the final the upper expected limit. . .	161
9.16	Differences (left) and pulls (right) for the measured and expected signal strengths for $m_{Z'} = 4000$, $m_{N_D} = 800$, $m_{N_S} = 200$. First row has an injected signal of 0, the second has the lower expected limit, the third the median expected limit, and the final the upper expected limit. . .	162
9.17	Differences (left) and pulls (right) for the measured and expected signal strengths for $m_{Z'} = 5500$, $m_{N_D} = 1800$, $m_{N_S} = 200$. First row has an injected signal of 0, the second has the lower expected limit, the third the median expected limit, and the final the upper expected limit. . .	163
9.18	Asymptotic fiducial cross section upper limits for 2016-2018 are shown in the m_{N_D} and $m_{Z'}$ plane for the $m_{N_S} = 200$ GeV scenario. Limits shown in fb.	165
9.19	Asymptotic fiducial cross section upper limits for 2016-2018 are shown in the m_{N_D} and $m_{Z'}$ plane for the $m_{N_S} = 1$ GeV scenario. Interpolated signals are not included. Limits shown in fb.	165
9.20	Asymptotic fiducial cross section upper limits for 2016-2018 as a function of m_{N_D} for fixed values of $m_{Z'}$ for the $m_{N_S} = 200$ GeV scenario. Limits shown in fb.	166
9.21	Asymptotic fiducial cross section upper limits for 2016-2018 as a function of m_{N_D} for fixed values of $m_{Z'}$ for the $m_{N_S} = 1$ GeV scenario. Limits shown in fb.	167
10.1	Left: Latest CMS Experiment narrow-resonance dijet limits [36]. Right: Latest ATLAS Experiment narrow-resonance dijet limits on excited quark q^* models [92]. Other signal interpretations have similar limits. These results can be interpreted to set limits on the hadronic channel of a leptophobic Z'	169
A.1	Leading and subleading muon p_T in the signal region for a signal cross section of 10 fb. Boosted signal is possible, hence the addition of the last bin. No trigger scale factors applied.	182
A.2	Privately generated muon trigger scale factors for 2016 (top), 2017 (middle), and 2018 (bottom).	183

A.3	Muon ID scale factors applied. The top figure shows the luminosity weighted 2016 scale factors [84]. The middle shows 2017 [85], and the bottom shows 2018 [86].	184
A.4	Electron reconstruction scale factors applied. From top to bottom: 2016, 2017, 2018. [81].	186
A.5	Electron ID scale factors applied. From top to bottom: 2016, 2017, 2018. [80].	187
A.6	Left column: Generic btagging scale factors. Right column: Btagging scale factors for $t\bar{t}$ MC.	188
B.1	Z to $e\mu$ at preselection in the jet mass signal region featuring H candidate kinematic distributions.	189
B.2	Z to $e\mu$ at preselection in the jet mass signal region featuring dilepton fake Z candidate kinematic distributions.	190
B.3	Z to $e\mu$ at preselection in the jet mass signal region featuring the RJR mass estimators.	190
B.4	Z to $e\mu$ at preselection in the jet mass signal region featuring mass variables, and btagging values.	191
B.5	Z to $e\mu$ at preselection in the jet mass signal region kinematic for the leading lepton.	191
B.6	Z to $e\mu$ at preselection in the jet mass signal region kinematic for the subleading lepton.	192
B.7	Z to $e\mu$ at preselection in the jet mass signal region kinematics for the muons selected to serve as the background estimation for the $\mu\mu$ channel.	192
B.8	Z to $e\mu$ at preselection in the jet mass signal region kinematics for the electrons selected to serve as the background estimation for the $\mu\mu$ channel.	193
B.9	Z to $e\mu$ at preselection in the jet mass signal region ΔR plots for Higgs candidate jet and Z constituents.	193
B.10	Z to $e\mu$ at preselection in the jet mass signal region $\Delta\phi$ plots for boson candidates and MET.	194
C.1	Additional closure plots in the validation region. Same selection as preselection signal region except the p_T^{miss} cut is reversed.	195

D.1	Comparison between the unscaled (left) and scaled by 10 (right) pull distributions for the $m_{Z'} = 2000$, $m_{N_D} = 400$, $m_{N_S} = 200$ signal mass point at the median expected limit injected signal strength.	197
D.2	Comparison between the unscaled (left) and scaled by 10 (right) pull distributions for the $m_{Z'} = 4000$, $m_{N_D} = 800$, $m_{N_S} = 200$ signal mass point at the median expected limit injected signal strength. Slight excess at negative pull values gone in the scaled version.	197
D.3	Comparison between the unscaled (left) and scaled by 10 (right) pull distributions for the $m_{Z'} = 5500$, $m_{N_D} = 1800$, $m_{N_S} = 200$ signal mass point at 0.0 injected signal strength. Slight excess at negative pull values gone in the scaled version.	198

List of Tables

2.1	Fields that interact under the new $U(1)_B$ charge in the minimal anomaly model and their associated charges. $SU(3)_C$, $SU(2)_W$, and $U(1)_Y$ represent the SM groups, with the associated columns showing the strong, weak isospin, and weak hypercharges associated with each of the new particles.[38].	26
5.1	Iterative tracking targets. Naming loosely corresponds to the seeding method used in each iteration. [63].	77
5.2	Summary of the clustering parameters. The neighboring cells are defined either as those cells sharing an edge with the seed cell (for a total of four neighbors), or additionally sharing corners (for a total of eight neighbors). The ECAL endcap has an additional requirement on the transverse energy due to a non-uniformity in the noise dependent on θ . Quoted values are determined from optimizations based photon, jet, π^0 , and K_L^0 simulation.	77
5.3	Applied p_T^{miss} filters. Filters are applied to both simulation and data unless otherwise specified. Filters handle known detector or reconstruction effects. The HCAL Phase 1 upgrade should remove the need for the HCAL filters.	84
6.1	Luminosities and uncertainties in Run 2 Collision Data recorded by the CMS Experiment.	88
6.2	Trigger paths used in Run 2 Collision Data.	88
6.3	The background MC used in the analysis with the corresponding cross section $\times$ branching fraction.	90
6.4	The criteria for the “Loose” electron ID. [61]. Subscript “SC” refers to the supercluster associated with the electron. A detailed description of the input parameters is given in section 6.4.1.	96
6.5	Comparison of Muon ID definitions	98
6.6	Tight Jet ID Definitions.	98

6.7	Signal yields with luminosity scaling for 2017 and 2018. The cross section is normalized to 10 fb. Cuts with an asterisk do not have btagging or muon ID scale factors applied.	106
7.1	Reconstructed Z candidate yields with > 0 fat jets and after p_T^{miss} filter application in the 2018D SingleMuon data. Z ($e\mu$) and Z ($\ell\ell$) candidates are reconstructed for each event. This table shows the number of events with multiple Z candidates reconstructed, which are vetoed. The overlap region is negligible.	111
7.2	Background shapes for Z + jets MC simulation normalization. The three background templates' ($f_{Z+jets}(x)$, $f_{t\bar{t}}(x)$, and $f_{VV}(x)$) parameters are fixed after fits to the respective MC. The total function $f_{\text{total}}(x)$ with the fixed background parameters is fit to the data sidebands with the normalization A allowed to float.	119
8.1	Summary of the affects of data and MC statistics on the α ratio extrapolation.	134
8.2	Summary of the effects of Jet Energy Scale corrections up and down shifts. For background processes that only have a “rate” systematic, the inclusion of the shape was tested in the limit framework, and resulted in no change to the limit. For signal, it is clear from Fig. 8.10 that a shape treatment is not appropriate.	135
8.3	Summary of the effects of btagging up and down shifts. For signal, it is clear from Fig. 8.12 that a shape treatment is not appropriate, and the dominant behavior is captured in the histogram normalization.	137
8.4	Summary of the effects of unclustered MET up and down shifts. For signal, it is clear from Fig. 8.14 that a shape treatment is not appropriate, and the dominant behavior is captured in the histogram normalization.	139
8.5	$t\bar{t}$ estimation HEM effects on input yields.	144
8.6	Summary of the effects on signal region yield of pdf scale shifts. Signal effects get larger with Z' mass.	145
8.7	Summary of the effects on signal region yield of maximal qcd scale shifts. Signal effects get larger with Z' mass.	146
9.1	Signal Injection Test Parameters.	160

9.2	Limit progression applying the data-driven $t\bar{t}$ background estimation method.	164
A.1	Binning used in the efficiency measurements and scale factor generation.	182
A.2	Summary of the tag and probe selections.	182

Chapter 1

Introduction

The 2012 discovery of the Higgs Boson by the ATLAS and CMS experiments redefined the landscape of particle physics discovery [2, 3]. This seminal moment completed the standard model (SM) of particle physics, thereby solidifying our descriptive model of three of the fundamental forces and validating our model for the origin of the particle masses. Despite this unequivocal triumph, questions about the fundamental nature of the universe remain. The role of a high energy particle physicist is to unearth the fundamental construction of the universe in terms of the particles that make up matter and the forces that govern interactions. The SM has given clear predictions; but now, we are bereft of its predictive power. Without a clear path to address the limits of the SM, we are not just free – we are obligated – to look for all that is possible.

Between 1960 and 2012, there were clear candidates for the next particle discovery. Each of these confirm the SM as a framework for understanding the electromagnetic, weak, and strong interactions of the fundamental particles; however, a description of gravity is flagrantly omitted in the SM. The astronomical observation of dark matter relegates the percentage of the universe the SM describes to a woeful 5%. These are two stark indicators among many others that the SM is incomplete. The shortcomings of the SM must be accounted for in Beyond the standard model (BSM) physics.

The BSM searches of the Large Hadron Collider (LHC) physics era have primarily been concerned with the simplest extensions of the SM; however, over a decade of operation and analysis have yet to yield solid evidence for new physics through these minimal models. With the completion of the canonical SM and the absence of new signatures in flagship channels, we can now explore the weird and unexpected remaining space unencumbered by expectation. This is a distinct re-definition of our duty of the past sixty years.

New physics can be hiding at LHC energies in complex, and therefore unexplored, final states. We must push our detectors and analysis models to the limit. The ATLAS and CMS detectors are optimized for the discovery of the Higgs Boson. While they are also well-suited for the more challenging task of discovering the unknown, planned detector upgrades can bolster their power. Many next-to-minimal models and exotic extensions remain unexplored, and often present challenging topologies that demand innovation in analysis practice. This thesis emphasizes readiness for the unexpected

in both hardware and analysis design.

Before Run 3 of the LHC, the CMS Hadron calorimeter (HCAL) underwent the so-called “Phase 1 Upgrade.” In this upgrade, the on-detector electronics for detector control and read out were completely replaced to extend the lifetime of the detector and to increase the sensitivity. These upgrades are now in use in the ongoing Run 3 data-taking of the LHC. This thesis also presents the full-Run 2 search of a leptophobic Z' decaying via anomalous in the $ZH \rightarrow \mu\bar{\mu}b\bar{b}$ and p_T^{miss} final state. The leptophobic-anomalous model is new in CMS, and uses a cascade decay to dark matter to hide a Z' at LHC energies. The model’s resonant cascade is a striking signature, and employs the first use of Recursive Jigsaw Reconstruction [4] in CMS to build an observable directly sensitive to the mass of the Z' . This analysis is a proof of concept not only of the discovery power of novel observables, but of the virtue of exploring unorthodox models. The High Luminosity LHC (HL-LHC) is on the horizon, and promises unprecedented amounts of data for exploration. Run 2 and 3 of the LHC can prioritize curiosity and creativity in analysis design that emphasizes broadening physics reach and hones our analysis tools moving into high luminosity running.

The structure of this thesis is as follows. Chapter 2 introduces the SM and motivates the leptophobic Z' model that is the inspiration of the search. Chapter 3 introduces the LHC and the CMS Experiment, and Chapter 4 describes the author’s involvement in the CMS HCAL Phase 1 upgrade. After a solid foundation in the physical detection of particles and the operational practice of running an experiment, Chapter 5 and 6 cover the event reconstruction and analysis strategy, respectively. Of principle importance to the analysis philosophy are the data-driven background methods which are described in Chapter 7. Associated systematics uncertainties are presented in Chapter 8. Finally, Chapter 9 and 10 include the results and conclusions.

Chapter 2

Theoretical Underpinnings

2.1 The Standard Model of particle physics

The standard model (SM) of particle physics is a quantum field theory (QFT) that describes and predicts the behavior of fundamental particles. A QFT extends the tenets of quantum mechanics to the relativistic regime, and casts particles as excitations of their underlying fields. Section 2.1.1 reviews the particles of the SM, and their general characteristics. The SM interactions are described by the crossing of three symmetry groups,

$$SU(3) \times SU(2) \times U(1),$$

each (roughly) capturing the behavior one of the three fundamental forces. Out of the four fundamental forces: electromagnetism, the weak force, the strong force, and gravity; the standard model encompasses the first three. The SM's relationship with gravity is an open question, and will be discussed briefly in Section 2.1.3.

At the heart of each of these symmetries is a conserved quantity in nature that dictates the rules by which the particles of the visible universe interact. Local transformations under the symmetry, analogous to coordinate transformations in spatial symmetries, should not change the underlying interaction, and should maintain the conserved quantity. The stipulation that these symmetries remain intact under local phase transformations promotes these global symmetries to local symmetries, a requirement termed *gauge invariance*. This produces the force-mediating particles, or gauge bosons, of the SM. The interactions of the SM are detailed in Section 2.1.2.

2.1.1 Particles

The particles of the SM can be divided into two general groups: the fermions, with 1/2-integer spin, and the bosons, with integer spin. Following different transformation rules, the fermions make up the known matter present in the universe, and the bosons, with the exception of the Higgs Boson, carry the forces of the SM. Even though the particles of the standard model can be broadly classified as fermions and bosons, this




Figure 2.1: Summary of the SM particles and their properties. [5, 6, 7]

simplification disguises a rich diversity. Figure 2.1 summarizes the particles of the standard model.

The fermions of the SM are divided into leptons and quarks. Under the SM, all leptons interact via the weak force, and charged leptons also interact via the electromagnetic force. Quarks participate in both weak and electromagnetic interactions, and additionally interact through the strong force. Both types of fermions exist in three generations, characterized by increasing charged-lepton and quark mass. Everyday matter consists mainly of the first generation of particles, while the second and third generations are accessible at colliders, like the Large Hadron Collider, discussed in Chapter 3. All spin-1/2 particles are described by the Dirac Lagrangian density,

$$\mathcal{L}_{Dirac} = \bar{\psi}(i\gamma^\mu\partial_\mu - m)\psi, \quad (2.1)$$

where γ^μ are the Dirac matrices, and ψ are the spin-1/2 fields.

The three lepton generations consist of an electrically charged lepton, the *electron*, *muon*, or *tau lepton*, and an associated neutral lepton, the *neutrino*. The charged leptons each carry an electric charge of -1, in units of positron charge. The neutrinos are nearly massless, and are treated as such in the realm of collider physics.

The three quark generations each contain an up- and down-type quark, with a fractional electric charge of $+2/3$ and $-1/3$, respectively, in units of positron charge. For the three generations, the up-type quarks are the *up*, *charm*, and *top*, and the down-type quarks are the *down*, *strange*, and *bottom* (sometimes called *beauty*). Like the electric charge in electromagnetism, color is the conserved charge in the strong

interaction. Under color interactions, generally speaking, quarks must exist in colorless bound-states of integer electric charge, creating composite objects known as hadrons. This concept is known as *asymptotic freedom*, and is covered in more depth in Section 2.1.2.

In the SM, the bosons can be grouped into spin-1 gauge, or vector, bosons, and the spin-0 scalar Higgs. The photon, gluon, Z and W are all gauge bosons, since they arise naturally to preserve the gauge-invariance of the SM Lagrangian. The gauge bosons are further classified as massive (the Z and W bosons) and the massless (photons and gluons), depending on the nature of their underlying symmetry. The Z and W bosons mediate the weak interaction, the photons the electromagnetic, and gluons the strong. Photons, gluons, and Z bosons are electrically neutral, while the W carries an electric charge to mediate the charged-current weak interaction. Gluons carry a color charge. The Higgs boson is a massive, electrically neutral, and colorless scalar field that is responsible for the masses of the other fundamental particles. Each of these forces and the role of the Higgs boson are covered in more detail in Section 2.1.2.

This discussion has completely omitted the inclusion of anti-particles: particles with the same mass as their particle counterparts, but with opposite-values of all quantum numbers except spin. The transition between particles and anti-particles is known as *charge conjugation*. Most striking is the inversion of the electric charge for charged particles. The anti-particle of a spin-up electron is a spin-up positron. Antimatter is captured naturally as the negative-energy solutions to the Dirac equation (Eq. 2.1) for fermions. Once one interprets these solutions as positive-energy states of opposite-sign, same mass particles, they take on physical meaning, as these states are observed in nature. Of the fundamental particles, only the photon and Z^0 is a known-eigenstate of the charge conjugation operator, meaning that they are their own anti-particles.

2.1.2 Interactions

Each of the fundamental forces that take part in SM interactions can be described by gauge theories. The Maxwell equations can be written in terms of a three-vector potential $\mathbf{A}$ that leaves the resulting fields invariant under the arbitrary gauge transformation

$$\mathbf{A} \rightarrow \mathbf{A} + \nabla\Lambda,$$

where Λ is an arbitrary scalar. In the Lagrangian formalism, this is associated with the conserved electric charge, and is therefore a consequence of a symmetry in the interaction. If one were not to know of the vector potential and its associated field energy a priori, one can arrive at the same form purely from the *requirement* of gauge invariance, as discussed in Section 2.1.2.1. In the quantized QFT, the vector

potential becomes the photon, the gauge boson of electromagnetism. When applying the parton model to understand the strong force and its relationship between the plethora of quark bound states, it was noticed that the partons seemed to transform under an $SU(3)$ symmetry. As a result, the gauge invariance machinery postulated the gluon fields prior to their discovery. The charged-current of the weak interaction, mediated by W bosons, can be captured as an $SU(2)$ symmetry; however, this leaves out the Z boson, and does not include the masses of the weak vector bosons in a gauge invariant way.

While it is natural to think of the three forces separately, one of the triumphs of the SM is the unification of the weak and electromagnetic forces into the Electroweak (EW) interaction, and the provision of weak boson masses by Electroweak Symmetry Breaking (EWSB). This unified description exploits the gauge invariant behavior of the $SU(2) \times U(1)$ group to incorporate the Z boson, and breaks that symmetry via the Higgs Mechanism to provide masses to the weak-interaction force mediators. The strong interaction eludes attempts at unification and still exists separately. The following discussion will begin with a description of electromagnetism in Section 2.1.2.1, and end with electroweak symmetry breaking in Section 2.1.2.4. Electromagnetism will be discussed twice, once to introduce the concept of a gauge boson, and a second time in Electroweak unification. The strong force is described in Section 2.1.2.2 as a generalization of the mathematics of the electromagnetic interaction to higher dimensional symmetry groups. The following discussion mainly follows Refs [8, 9].

2.1.2.1 Electromagnetic Interaction

The electromagnetic interaction is described by Quantum Electrodynamics (QED). Quantum Electrodynamics captures the interactions between electrically charged particles via the exchange of the photon. The central tenet of QED is the conservation of electric charge, implying a $U(1)$ symmetry in the theory. To first order, photons do not self-interact, so QED can be described by an Abelian gauge theory.

The QED Lagrangian Density is given by

$$\mathcal{L}_{QED} = \bar{\psi}i\gamma^\mu(\partial_\mu - m)\psi - \frac{1}{4}(F_{\mu\nu})^2 - Q\bar{\psi}\gamma^\mu\psi A_\mu, \quad (2.2)$$

where ψ is a fermion field, A_μ is the electromagnetic vector potential (photon), $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the electromagnetic field tensor, and Q is the electric charge. The first term is simply the Dirac Lagrangian, the second Maxwell's, and the third term represents the interaction between the fermion field and the photon.

This Lagrangian can also be made simply by requiring that the Dirac Lagrangian given by Eq. 2.1 is invariant under the $U(1)$ transformation

$$\psi(x) \rightarrow e^{i\alpha(x)}\psi(x). \quad (2.3)$$

Inner product terms with the $m\bar{\psi}\psi$ are invariant, but the derivative terms are more complicated. To capture the gauge invariant behavior, the *gauge covariant derivative* D_μ in the arbitrary direction n^μ is defined as

$$n^\mu D_\mu \psi = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} [\psi(x + \epsilon n) - U(x + \epsilon n, x)\psi(x)], \quad (2.4)$$

where

$$U(y, x) \rightarrow e^{i\alpha(y)}U(y, x)e^{-i\alpha(x)}. \quad (2.5)$$

Since $U(y, x)$ should be smooth and well-behaved, $U(x + \epsilon n, x)$ can be expanded as

$$U(x + \epsilon n, x) = 1 - iQ\epsilon n^\mu A_\mu(x) + \mathcal{O}(\epsilon^2), \quad (2.6)$$

where Q is an arbitrary constant, and A_μ is a new vector field. The covariant derivative in Eq. 2.4 then becomes

$$D_\mu = \partial_\mu + iQA_\mu. \quad (2.7)$$

With this new expression replacing the partial derivative in Eq. 2.1, A_μ is introduced into the Lagrangian as a new vector field. To preserve the gauge invariance, both D_μ and A_μ must also be invariant under the transformation in Eq. 2.3. The transformation rules of the covariant derivative and the vector field A_μ help define the final, missing piece of the QED Lagrangian: the kinetic terms of A_μ . The invariance of the covariant derivative under the same transformation of the fermion field implies that the commutator of the covariant derivative is also invariant under

$$[D_\mu, D_\nu]\psi(x) \rightarrow e^{\alpha(x)}[D_\mu, D_\nu]\psi. \quad (2.8)$$

Expanding the commutator gives

$$\begin{aligned} [D_\mu, D_\nu] &= iQ(\partial_\mu A_\nu + A_\mu \partial_\nu - \partial_\nu A_\mu - A_\nu \partial_\mu), \\ &= iQ(\partial_\mu A_\nu - \partial_\nu A_\mu). \end{aligned} \quad (2.9)$$

The last line of Eq. 2.9 produces the familiar electromagnetic field tensor,

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu, \quad (2.10)$$

of the QED Lagrangian given in Eq. 2.2. From the gauge invariant perspective, the complete form of the QED Lagrangian emerges. Written expressly in gauge invariant form, the QED Lagrangian is

$$\mathcal{L}_{QED} = \bar{\psi} i \gamma^\mu (D_\mu - m) \psi - \frac{1}{4} (F_{\mu\nu})^2.$$

2.1.2.2 Strong Interaction

The strong interaction consists of the interactions of quarks with *color* charge, and is described by Quantum Chromodynamics (QCD). There are three color charges, nominally termed *red*, *blue*, and *green*, therefore the colored quark interactions are captured as transformations of the $SU(3)$ group. Early experiments revealed that the quantum field theory describing quark interactions must be *asymptotically free*. Asymptotic freedom is the behavior where as the momentum of interacting particles increases, (at a collider, for example), the coupling-strength of the interaction decreases. This property explains first why no free-quarks were observed in nature, and why as the energy increases, free-particle descriptions of quarks can be used to derive experimental observables. This trend in coupling strength is a feature of non-Abelian gauge theories.

To this end, the generators of the $SU(3)$ of QCD follow

$$[T^a, T^b] = i f^{abc} T^c, \quad (2.11)$$

where $T^a = \lambda^a/2$, λ^a are the Gell-Mann matrices, and f^{abc} are the structure constants of $SU(3)$. The Gell-Mann matrices are a set of eight, three-by-three Hermitian and trace-less matrices.

To preserve gauge invariance, we stipulate that the theory is invariant under the transformation

$$q \rightarrow e^{ig_s T^a} q, \quad (2.12)$$

where q is the quark field, and g_s is the strong coupling constant. This parallels the QED case, but with the generator of the group providing the phase. Since quarks are fermions, this new interaction must preserve the gauge invariance of the Dirac Lagrangian given in Eq. 2.1, and the covariant derivative in this case is redefined to be

$$D_\mu = \partial_\mu + ig_s T^a G_\mu^a, \quad (2.13)$$

producing the eight gluon fields G_μ^a as the strong force mediators. The gluons add new kinetic terms to the Lagrangian density, which can be found using the transformations of the covariant derivative.

The covariant derivative is expected to be invariant under the same transformation of quark field, which in turn implies that the commutator of the covariant derivative is also invariant under

$$[D_\mu, D_\nu]q \rightarrow e^{ig_s T^a} [D_\mu, D_\nu]q. \quad (2.14)$$

In the same fashion as the derivation of the electromagnetic field tensor in Eq. 2.9, the commutator of the QCD covariant derivative can be expanded to produce the field tensor of the gluons. Since the generators of the $SU(3)$ are non-Abelian, the above commutation relation will naturally bring cubic and quartic interaction terms among the eight gluons. This results in the gluon field tensor

$$G_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a - g_s f_{abc} G_\mu^b G_\nu^c. \quad (2.15)$$

With the propagation of the gluon described, the QCD Lagrangian becomes

$$\mathcal{L}_{QCD} = \bar{q} i \gamma^\mu D_\mu q - \frac{1}{4} G_{\mu\nu}^a G^{\mu\nu a}, \quad (2.16)$$

with an implicit summation over quark flavor, which has been ignored thus-far.

2.1.2.3 Electroweak Interaction

While QED can exist on its own as a renormalizable, Abelian gauge theory, the weak interaction with three vector bosons, the W^+ , W^- , and Z , cannot be captured in a stand-alone gauge theory satisfactorily. As it turns out, the three massive vector bosons of the weak interaction and the single, massless photon of QED must be treated together as an $SU(2) \times U(1)$ symmetry to arrive at the weak force observed in the natural world.

Contrary to the nature of electromagnetism, the charged interaction of the weak force is parity-violating, therefore the description of the weak force must be chiral in nature, and differentiate between “left-handed” and “right-handed” particles. The charged-current weak interaction exclusively interacts with the left-handed projections of particles. The reverse is true for anti-particles, where only the right-handed

projections interact. Electromagnetism and the weak neutral current make no such distinction, and a unified theory must respect both of these transformation properties.

Projection operators P_L and P_R produce the left and right-handed components of a fermion field ψ , respectively

$$\begin{aligned}\psi_L &= P_L \psi = \frac{1}{2}(1 - \gamma_5)\psi, \\ \psi_R &= P_R \psi = \frac{1}{2}(1 + \gamma_5)\psi,\end{aligned}\tag{2.17}$$

where γ_5 is defined as $\gamma_5 = i\gamma_0\gamma_1\gamma_2\gamma_3\gamma_4$. For charged-current interactions, the left-handed field components of same-generation leptons and quarks transform as doublets under $SU(2)$, while the right-handed components transform as singlets, negating any interaction. This preserves the chiral structure.

Both of the symmetry groups in the $SU(2) \times U(1)$ of the electroweak interaction have associated quantum numbers. The *weak isospin*, T is the quantum number associated with the $SU(2)$, and takes the eigenvalues of the associated three generators, $T_i = \frac{\sigma_i}{2}$, where σ_i are the three Pauli matrices. The quantum number of the $U(1)$ is the *weak hypercharge*, Y , and can be assigned to ensure that QED interactions are recovered. The electroweak quantum numbers are related to the conserved electric charge via

$$Q = T + Y.\tag{2.18}$$

Using the knowledge of the gauge structure of pure QED and QCD, the covariant derivative of the EW interaction can be written as

$$D_\mu = \partial_\mu - ig_w T^i W_\mu^i - ig_{\text{em}} Y B_\mu,\tag{2.19}$$

where the W_μ^i are the three gauge boson fields of the $SU(2)$, and B_μ is the boson field associated with the EW $U(1)$. In addition to the weak hypercharge Y , the above covariant derivative also includes two separate couplings, g_w and g_{em} , for the two separate symmetry groups. Following a similar invariance transform as Eq. 2.14, the Lagrangian of these new gauge bosons is given as

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}W_{\mu\nu}^i W^{i\mu\nu} - \frac{1}{4}B_{\mu\nu} B^{\mu\nu},\tag{2.20}$$

where

$$W_{\mu\nu}^i W^{\mu\nu i} = \partial_\mu W_\nu^i - \partial_\nu W_\mu^i + g_w \epsilon^{ijk} W_\mu^j W_\nu^k,\tag{2.21}$$

and

$$B_{\mu\nu}B^{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu. \quad (2.22)$$

All together, taking the pair of the electron and the electron neutrino fields (e and ν_e) as an example, one can write the massless (gauge invariant) electroweak Lagrangian as

$$\begin{aligned} \mathcal{L}_{EW} = & i\bar{L}\gamma^\mu D_\mu L + i\bar{e}_R\gamma^\mu(\partial_\mu - ig_{em}Y_R^{(e)}B_\mu)e_R \\ & + i\bar{\nu}_R\gamma^\mu(\partial_\mu - ig_wY_R^{(\nu)}B_\mu)\nu_R + \mathcal{L}_{gauge}, \end{aligned} \quad (2.23)$$

where L is the doublet of left-handed electron and neutrino projections,

$$L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix} = \begin{pmatrix} \frac{1}{2}(1 - \gamma_5)\nu \\ \frac{1}{2}(1 - \gamma_5)e \end{pmatrix}, \quad (2.24)$$

e_R and ν_R are the right-handed projections (which transform as singlets)

$$e_R = \frac{1}{2}(1 + \gamma_5)e, \quad \nu_R = \frac{1}{2}(1 + \gamma_5)\nu, \quad (2.25)$$

and $\mathcal{L}_{gauge}$ is defined in Eq. 2.20. The right-handed neutrino projection can be ignored. In this Lagrangian, the correct number of bosons exist, but the true SM bosons are linear combinations of the above gauge fields. The SM $W^\pm$ are built from the linear combination

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2). \quad (2.26)$$

This now leaves the W_μ^3 and B_μ fields as the candidates for the Z and the photon. Ignoring the components of Eq. 2.23 that involve the $W^\pm$, one is left with the Lagrangian

$$\begin{aligned} \mathcal{L}_{NC,EM} = & \frac{1}{2}g_w\bar{\nu}_L\gamma^\mu\nu_L W_\mu^3 - \frac{1}{2}g_w\bar{e}_L\gamma^\mu e_L W_\mu^3 + g_{em}Y_L\bar{n}u_L\gamma^\mu\nu_L B_\mu + \\ & g_{em}Y_L\bar{e}_L\gamma^\mu e_L B_\mu + g_{em}Y_R^{(e)}\bar{e}_R\gamma^\mu e_R B_\mu + g_{em}Y_R^{(\nu)}\bar{\nu}_R\gamma^\mu\nu_R B_\mu. \end{aligned} \quad (2.27)$$

By inspection, neither of the gauge fields in Eq. 2.27 can be the Z or the photon directly, since no values of Y give their respective behaviors. Both fields couple to

the neutrino, which is not the case of the QED photon. The Z and A_μ (photon) can be related to the W_μ^3 and B_μ as

$$\begin{aligned} Z_\mu &= \cos \theta_W W_\mu^3 - \sin \theta_W B_\mu, \text{ and} \\ A_\mu &= \sin \theta_W W_\mu^3 + \cos \theta_W B_\mu, \end{aligned} \quad (2.28)$$

where θ_W is known as the *weak mixing* or “Weinberg” angle. This orthogonal transform retains the ability to write the kinetic gauge boson terms of the Lagrangian without any mixing. With the vector fields decomposed in this manner, the weak hypercharges Y can be assigned values to give the appropriate behavior of QED. With these physical constraints, the value of θ_W can be found to be

$$e = g_w \sin \theta_W, \quad (2.29)$$

where e in this case has returned to be the value of the electric charge. This discussion presents a mathematical justification of the weak-vector bosons as gauge bosons, illustrated through their couplings to first-generation leptons. Phenomenologically, the quark sector also carries a weak charge and interacts with the weak vector bosons. Quark interactions via the charged-weak current are chiral, just as in the lepton case; however, flavor need not be conserved in these interactions. To account for this, the quark doublets that interact under the $SU(2)$ symmetry of the charged-current weak interaction are defined as

$$\begin{pmatrix} u \\ d' \end{pmatrix}, \quad \begin{pmatrix} c \\ s' \end{pmatrix}, \quad \begin{pmatrix} t \\ b' \end{pmatrix}, \quad (2.30)$$

where the “primed” values are given by

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \quad (2.31)$$

In the above relation, the matrix V is the *Cabbibo-Kobayashi-Maskawa* (CKM) matrix, and the elements represent the coupling of the $j \rightarrow i + W^-$ transition, V_{ij} [10, 11]. The CKM matrix represents the relationship between the mass eigenstates of the down-type quarks and their weak-interaction doublets. While the diagonal values are close to one, the off-diagonal elements represent the generational mixing present in the weak force.

Up until now, there has been no mention of the masses of the weak-force mediating bosons. One of the original hints that unification was appropriate was that for all

intents and purposes, electromagnetism and the weak force are similar, but the range and strength of the fields differed. This is explained by the fact the W and Z have mass, but the origin of this mass remained a mystery, and could not be included in a gauge invariant way. For a “unified theory,” this is rather unbecoming. Beyond the gauge boson sector, the mass terms present in the fermion Lagrangians are not motivated. While they can be included without violating gauge invariance, the arbitrary m has no origin, and is simply included by hand.

In theories with intact symmetries, like QED and QCD, the gauge bosons are massless; however, masses can be included into the Lagrangians in a gauge-invariant way if the symmetry is broken. Therefore, the symmetry spawning the weak-force mediators must be a broken symmetry. The Higgs Mechanism simultaneously breaks the symmetry of the electroweak $SU(2) \times U(1)$ group, and illustrates how the addition of a scalar field maintains the gauge invariance of the theory.

2.1.2.4 Higgs Mechanism

The first step to understanding the Higgs Mechanism [12, 13, 14] is to incorporate the spontaneous symmetry breaking into the model. This can be simply illustrated taking the linear sigma model [8] described by the Lagrangian

$$\mathcal{L} = \frac{1}{2}(\partial_\mu \phi^i)^2 + \frac{1}{2}\mu^2(\phi^i)^2 - \frac{\lambda}{4}[(\phi^i)^2]^2, \quad (2.32)$$

where ϕ^i are a set of N real scalar fields, the value μ is a real parameter with the dimension of mass, and the λ is a dimensionless, positive coupling constant. The Lagrangian is invariant under the symmetry transformation

$$\phi^i \rightarrow R^{ij} \phi^j$$

for any $N \times N$ orthogonal matrix R . To get more traction describing this theory, it is convenient to consider descriptions around the energy density minima for the potential term given by

$$V(\phi^i) = -\frac{1}{2}\mu^2(\phi^i)^2 + \frac{\lambda}{4}[(\phi^i)^2]^2. \quad (2.33)$$

Figure 2.2 visualizes Eq. 2.33 when $N = 2$. The minima correspond to any value of ϕ_0^i that satisfies

$$(\phi_0^i)^2 = \frac{\mu^2}{\lambda} = v^2, \quad (2.34)$$




Figure 2.2: Potential for spontaneously breaking a $N = 2$ symmetry. Minima would correspond to the vacuum expectation value. [8]

where v is the *vacuum expectation value* (VEV). These are naturally degenerate solutions. One can choose any direction, and therefore can choose the solution

$$\phi_0^i = (0, 0, \dots, 0, v). \quad (2.35)$$

Now that a direction has been chosen, the symmetry is *broken*. It is convenient to now consider the theory around the minimum, so the field $\phi^i(x)$ can be shifted to be

$$\phi^i(x) = (\pi^k, v + \sigma(x)), \quad k = 1, \dots, N - 1. \quad (2.36)$$

where π^k and σ are new spin-zero scalars. With this shift, the Lagrangian in Eq. 2.32 becomes

$$\begin{aligned} \mathcal{L} = & \frac{1}{2}(\partial_\mu \pi^k)^2 + \frac{1}{2}(\partial_\mu \sigma)^2 - \frac{1}{2}(2\mu^2)\sigma^2 \\ & - \sqrt{\lambda}\mu\sigma^3 - \sqrt{\lambda}\mu(\pi^k)^2\sigma + \text{higher orders}. \end{aligned} \quad (2.37)$$

This Lagrangian includes a massive new scalar σ and a set of massless scalar bosons, π^k . These new $N - 1$ massless bosons, called *Goldstone Bosons*, hide the original $O(N)$ symmetry.

The Higgs Mechanism is miraculous in that it takes the idea of Goldstone bosons and marries them to the idea of gauged boson fields. In this paradigm, the additional

scalar fields must be gauge invariant under the interaction's covariant derivative. Since these new scalars must couple to the $SU(2)$ of EW theory, it is convenient to define a complex doublet ϕ , comprised of four real scalar fields ϕ_i ,

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}, \quad (2.38)$$

where the $(+, 0)$ superscripts in Eq. 2.38 indicate weak isospin charges of $+1$ and 0 , respectively. The Lagrangian for this new field is given as

$$\mathcal{L}_{Higgs} = (D_\mu \phi^\dagger)(D^\mu \phi) - V(\phi), \quad (2.39)$$

where D_μ is the covariant derivative from Eq. 2.19 and the potential V is

$$V(\phi) = -\mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2. \quad (2.40)$$

Similarly to the purely Goldstone formulation, to minimize the potential, one can say that

$$\phi_0^\dagger \phi_0 = \frac{\mu^2}{2\lambda} = \frac{v^2}{2}, \quad (2.41)$$

where v is the vacuum expectation value. This product is an infinite set of doublets, and once one chooses a specific ground state, the symmetry is broken. It is useful to write ϕ in its phase representation,

$$\phi(x) = \exp\left(\frac{i}{v} \pi^a(x) \sigma^a\right) \rho, \quad (2.42)$$

where $\pi^a, a = 1, 2, 3$ are the Goldstone bosons, σ^a are the Pauli matrices, and

$$\rho^2 = \phi^\dagger \phi.$$

Just as in equations 2.35 and 2.36, one can choose a direction for ϕ , and shift to the minimum, yielding

$$\phi(x) = \exp\left(\frac{i}{v} \pi^a(x) \sigma^a\right) \begin{pmatrix} 0 \\ \frac{1}{\sqrt{2}}(v + H(x)) \end{pmatrix}, \quad (2.43)$$

where $H(x)$ is the conspicuously named additional scalar field, analogous to σ in Eq. 2.37. This has followed the same procedures as when the expressions for the

Goldstone bosons were derived. For a complete theory, the gauge invariance must be respected. Looking at Eq. 2.43, and recalling that gauge invariance implies that participating fields must be constant under an arbitrary phase change, one is free to choose to a gauge where the phase contribution is zero. This gauge is termed the *Unitary gauge*, and is equivalent to setting the remaining Goldstone Boson fields π^a to zero. In this way, the unphysical fields are removed from the Lagrangian, leaving

$$\phi(x)_U = \begin{pmatrix} 0 \\ \frac{1}{\sqrt{2}}(v + H(x)) \end{pmatrix}. \quad (2.44)$$

Plugging Eq. 2.44 into the Lagrangian from Eq. 2.39, making some convenient shifts, and doing some algebra yields

$$\begin{aligned} \mathcal{L}_{Higgs} = & \frac{1}{2} \partial_\mu H \partial^\mu H - \lambda v^2 H^2 - \lambda v H^3 - \frac{1}{4} \lambda H^4 \\ & + \frac{1}{8} (v + H)^2 (g_w^2 W_\mu^i W^{i\mu} - 4Y g_w g_{em} A_\mu^3 B^\mu + 4Y^2 g_{em}^2 B_\mu B^\mu). \end{aligned} \quad (2.45)$$

At this point, the additional scalar boson H has become massive, and takes the role of the Higgs Boson, interacting with the SM vector bosons. While the previously defined expressions for Z and the photon in Eq. 2.28 are not immediately obvious, it is clear that the quadratic terms in the Lagrangian 2.45 lead to the Z and W vector bosons acquiring mass. Taking the terms above that only involve the vector bosons, one can write

$$\begin{aligned} \mathcal{L}_{mass,ZW} = & \frac{1}{8} (g_w^2 + 4Y^2 g_{em}^2) v^2 \left(\frac{g_w}{\sqrt{g_w^2 + 4Y^2 g_{em}^2}} W_\mu^3 - \frac{2Y g_w}{\sqrt{g_w^2 + 4Y^2 g_{em}^2}} \right)^2 \\ & + \frac{1}{4} g_w^2 v^2 W_\mu^- W^{+\mu}, \end{aligned} \quad (2.46)$$

where $W_\mu^\pm$ are the W boson fields. To have the previous, massless results of Eq. 2.29, one must set the weak hypercharge above equal to $Y = 1/2$. This gives

$$Z_\mu = \frac{1}{\sqrt{g_w^2 + g_{em}^2}} (g_w W_\mu^3 - g_{em} B_\mu). \quad (2.47)$$

From Eq. 2.28, the photon is given orthogonally as

$$A_\mu = \frac{1}{\sqrt{g_w^2 + g_{em}^2}} (g_w W_\mu^3 + g_{em} B_\mu). \quad (2.48)$$

Since Eq. 2.48 does not appear in the mass terms of Eq. 2.46, it is clear that the photon is not given mass by the Higgs Mechanism. This is the desired result. All in all, for the Electroweak sector, the Higgs Mechanism produces boson masses

$$\begin{aligned} m_W &= \frac{1}{2}g_w v, \\ m_Z &= \frac{1}{2}(g_w^2 + g_{em}^2)^{1/2}v, \\ m_H &= v\sqrt{\lambda}. \end{aligned} \tag{2.49}$$

2.1.2.5 Fermion Masses

This Higgs field is also responsible for the masses of the fermion content of the SM. Mass terms can be added in a renormalizable manner via a *Yukawa* coupling, which has the form

$$\mathcal{L}_{Yukawa, generic} = \mathcal{L}_\psi + \mathcal{L}_\phi - g\bar{\psi}\psi\phi. \tag{2.50}$$

The chiral projections of the leptons must have the same mass, but the $SU(2) \times U(1)$ symmetry must be respected. Taking the electron and electron neutrino generation as an example, without loss of generality, this results in an interaction of the form

$$\mathcal{L}_{Yukawa} = -g_e \bar{L}\phi e_R + h.c. \tag{2.51}$$

where L has the same content (but a row structure) as Eq. 2.24, e_R is the right-handed electron singlet, g_e is a dimensionless coupling constant, $h.c.$ is the Hermitian conjugate, and the negative is added for convenience. In the unitary gauge, this becomes

$$\begin{aligned} \mathcal{L}_{Yukawa}^U &= -g_e(\bar{\nu}_L, \bar{e}_L) \begin{pmatrix} 0 \\ \frac{1}{\sqrt{2}}(v + H) \end{pmatrix} e_R + h.c., \\ &= -\frac{1}{\sqrt{2}}g_e(v + H)(\bar{e}_L e_R + \bar{e}_R e_L), \\ &= -\frac{1}{\sqrt{2}}g_e v \bar{e}e - \frac{1}{\sqrt{2}}g_e \bar{e}e H, \end{aligned} \tag{2.52}$$

producing an electron mass

$$m_e = \frac{1}{\sqrt{2}}g_e v. \tag{2.53}$$

This process can be repeated generally for the muon and the tau. The quarks of course can also acquire mass via Yukawa couplings to the Higgs, but care must be taken to account for the flavor-mixed structure of the $SU(2)$ doublets in Eq. 2.30. In spite of this, the mass relation ends up being the same, with the quark masses given as

$$m_q = \frac{1}{\sqrt{2}} g_q v. \quad (2.54)$$

While the quark-flavor mixing of the CKM matrix (Eq. 2.31) was first included due to experimental evidence, the Higgs mechanism naturally brings about the mixing. If one starts with left-handed doublets not necessarily corresponding to the physical fields, but to the generations, and proceeds to add Yukawa terms, the physical fields fall out as mass-eigenstates, but the original fields are of course superpositions of these. By generically implementing quark fields with EWSB in mind, the mixing appears without an ad hoc implementation.

The above discussion highlights the descriptive power of the SM, and omits a rigorous overview of the immense experimental verification of the predictive power of the SM. Figure 2.3 shows the remarkable agreement between the most recent Compact Muon Solenoid Experiment measurements of the EW interactions and their SM predictions. By accommodating most of the observed world of particle physics and presenting a relatively graceful explanation of three of the fundamental forces as the exchange of gauge bosons via the Higgs Mechanism, the SM appears to be complete.

Now we are back in the wild-west: there is no obvious candidate for the next discovery. No missing quark partner, as in the top quark. No missing scalar from symmetry breaking, as with the Higgs. That does not mean that the search should end.

2.1.3 Limitations and Open Questions

The Standard Model falls short of its primary function as a complete description of the particles in our universe. Generally speaking, there are two ways the SM fails as an underlying theory. First, the SM does not include all known objects and forces in the universe. One of the most glaring omissions is the force with humans are most familiar: gravity. The second is more subtle. We not only desire the SM to *describe* the universe, but we would also like it to *explain* observed phenomena. The Higgs boson *explains* the presence of the otherwise unmotivated fermion mass-terms in the SM Lagrangian. Before, they were added by hand with experimentally measured values. Like the fermion masses before EWSB, the SM does not explain all aspects of the observed universe, like the differences in the coupling strengths of the four fundamental forces.




Figure 2.3: Summary of current public SM EW interaction cross section measurements by the Compact Muon Solenoid Experiment. [15]

In the following subsections, a few of the open issues in the SM will be presented. Gravity (Section 2.1.3.1), Dark Matter (Section 2.1.3.2), and neutrino mass (Section 2.1.3.3) all provide examples of the first type of failing of the SM: the missing description of particles or interactions. The Hierarchy Problem (Section 2.1.3.4), and the pursuit of a Grand Unified Theory (GUT) (Section 2.1.3.5) seek to explain the nature of the SM.

2.1.3.1 Gravity

While it is obvious particles interact via gravity, the force carrier of gravity, termed the *graviton*, thwarts inclusion into the SM. When one attempts to quantize Einstein's equations for gravity, one is left with a non-renormalizable theory. Renormalization has thus far been a fundamental tenet of model building, since it avoids divergences in scattering amplitudes due to higher order corrections. Without renormalization, gravity can only be quantized at the Planck scale where the coupling of the gravitational interaction of particles is of order 1.

2.1.3.2 Dark Matter

Cosmological observations of galaxy rotation curves [16, 17], gravitational lensing [Massey_2010], and the cosmic microwave background early surveys of the universe [18] all reveal the presence of large amounts of “dark matter”, or matter that does not interact electromagnetically, in the universe. In fact, the SM only accounts for roughly five percent of the measured energy in the universe [18]. Dark matter offers a new class of particles that are not currently included in the SM. Measurements only reveal that Dark Matter interacts gravitationally, and that it is stable.

2.1.3.3 Neutrino Masses

Neutrinos are massless in the SM. Super-Kamiokande’s observance of neutrino-flavor oscillations [19] indicate that neutrinos must have mass, albeit small – in contrast to the SM expectation. Flavor oscillations arise if the neutrino flavor eigenstates are not the same as the neutrino mass eigenstates. The Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [10.1143/PTP.28.870, 20] maps the flavor-mass mixing, but the origin of the mixing (and mass) is not accounted for in the SM.

2.1.3.4 Hierarchy Problem

Since the SM is inherently incomplete, it is natural to suppose that the SM is an “effective field theory,” or a QFT that is the low-energy approximation of some higher-energy theory. One would expect then that the vacuum expectation value present in EWSB, and therefore the particle masses, is a component of that larger field theory. The fundamental, high-energy theory should also include gravity. QFTs are renormalized up to some scale Λ_{cutoff} , but in the case of some larger-energy theory,

$$\Lambda_{\text{cutoff}} \rightarrow \Lambda_{\text{Planck}} ,$$

where Λ_{Planck} is the Planck scale. The Planck scale is given as

$$\Lambda_{\text{Planck}} \rightarrow m_{\text{Planck}} = \left(\frac{G_N}{\hbar c}\right)^{-1/2} 10^{19} \text{GeV} . \quad (2.55)$$

The Planck scale is of about sixteen orders of magnitude above the EW scale, and this is the nature of the Hierarchy problem: the large discrepancy in the fundamental scales of the universe, and how these scales should impact observed phenomena.

In scalar QFTs like the for Higgs, the first order loop corrections for all interacting fermion terms like $-\lambda_\psi H \bar{\psi} \psi$ contribute to the $(\text{mass})^2$ term as [21]

$$\Delta m_H^2 = -\frac{|\lambda_\psi|^2}{8\pi^2} \Lambda_{\text{cutoff}} + \dots \quad (2.56)$$

When $\Lambda_{\text{cutoff}} \rightarrow \Lambda_{\text{Planck}}$, these contributions become very large, blowing up the resulting Higgs Mass. This is incompatible with the experimental observation of $m_H = 125$. This discrepancy is known as the “hierarchy” or “naturalness” problem.

2.1.3.5 Grand Unified Theories

The unification of the electromagnetic and weak forces sets a tantalizing precedent that the use of larger symmetry groups with a broken symmetry can unify all of the forces of the SM. The running of the coupling constants suggests that a Grand Unified Theory (GUT) is possible. The simplest group possible for this unification is the $SU(5)$ symmetry group, though it has been largely excluded by experiment [22]. The quest for larger symmetry groups remains, since the GUT symmetry breaking scale is close (in the way unreasonable numbers can be considered close) to the Plank scale.

2.2 Heavy vector bosons: A Common Beyond the Standard Model Ingredient

While the SM has known shortcomings, any new model to address these issues must preserve the predictive and descriptive power of the SM at the electroweak scale. A feature of many beyond the standard model (BSM) scenarios is an additional $U(1)$ symmetry, resulting in the inclusion of a massive Z' . This thesis presents a search for such a particle.

A Z' naturally arises out of models developing Grand Unified Theories or addressing the Hierarchy problem. The high energy behavior inherent in both problems forces the broadening of the fundamental theory to include new particles and phenomena. Both classes of models involve larger symmetry groups, and these groups can be broken to the $SU(3) \times SU(2) \times U(1)$ of the SM. This breaking often results in an additional $U(1)$ group. The symmetry breaking and Z' arise analogously to the massive Z boson in EWSB described in Sections 2.1.2.3 and 2.1.2.4. Other symmetry breaking mechanisms exist, but do not change the fundamental correspondence between the massive Z' and the $U(1)$ group.

Experimental searches for new physics are best when motivated by models addressing one, and optimally several, of the issues with the SM summarized in Section 2.1.3.

For example, models with Large Extra Dimensions not only address the scale discrepancies of the Hierarchy Problem in Section 2.1.3.4, but also include a quantum description of gravity. Additional $U(1)$ symmetries further result in additional associated particles. These new particles can readily provide dark matter candidates, making Z' physics rich not only due to the fundamental structure of the BSM model, but also for the phenomenological implications.

Section 2.2.1 discusses the general features of select theoretical models that address the limitations presented in Section 2.1.3, and presents how these models result in a Z' . The mentioned models represent a generic but non-exhaustive list of BSM scenarios that result in Z' . With the motivation of a search for a Z' established, Section 2.2.2 introduces the benchmark Z' model for this analysis.

2.2.1 Overview of Heavy Vector Bosons

For GUTs from extended symmetry groups, the number of additional gauge bosons possible in the group G is given as $n = \text{rank}[G]$ [23]. Therefore, symmetry groups larger than $SU(5)$ have at least one additional neutral gauge boson. The smallest number of additional neutral gauge bosons is in $SO(10)$. $SO(10)$ models also have additional fermions that can be considered dark matter candidates [24]. The larger groups behave similarly, with increasing numbers of additional bosons and fermions. The next, higher rank GUT models are based on the E_6 group. The E_6 exceptional group is the only exceptional group that is consistently anomaly-free, and has complex representations [25]. Anomalies will be discussed in more detail in Section 2.2.2. This makes E_6 theories suitable for unification. The caveat to these models is that there is no mechanism that prevents the Z' from being at the Planck scale, potentially prohibiting their detection at current and prospective particle colliders [26].

Supersymmetry (SUSY) stabilizes the Higgs mass, addressing the Hierarchy problem described in Section 2.1.3.4. In SUSY, if for every SM boson there is a corresponding *sfermion* and for every fermion there is a *gaugino*, the corrective terms given in Eq. 2.56 are perfectly canceled, stabilizing the mass of the Higgs Boson [27]. This mixing of fermions and bosons was shown to be a symmetry with a definite algebra by J. Wess and B. Zumino in 1974 [28]. If SUSY were to be a perfect symmetry, the masses of the additional particles would match their SM counterparts. Since no SUSY particles have been observed, SUSY particles must have masses due to a broken symmetry, warranting higher rank symmetry groups. SUSY at the TeV scale can be captured via broken E_6 models, producing a Z' which can decay to sfermions [29]. While the symmetry group structure is the same as those used in GUT models, Supersymmetric conditions keep the Z' mass away from the Planck scale, and closer to those accessible at colliders [30].

Composite Higgs models modify the nature of the Higgs boson, re-framing the hi-

erarchy problem. In composite models, there are three fundamental scales in the interaction: the EW scale, the confinement scale of composite Higgs constituents, and the larger free-particle scale [31]. Treating each of these scales separately at renormalization removes the impact of the corrections to the Higgs mass. In minimal composite models, the representation of the “ultrafermions” that comprise the Higgs bound state is given as

$$G_{UC} \times SU(2) \times U(1), \quad (2.57)$$

where G_{UC} is the “ultracolor” gauge group corresponding to the confinement scenario of the ultrafermions [31]. At the scale where the ultrafermion condensate –the Higgs– forms, the electroweak breaking of $SU(2) \times U(1)$ proceeds with an additional $U(1)$ symmetry, producing the Z' .

Extra Dimensional Models (both Large and Small) address the Hierarchy problem by removing the necessity of renormalization up to the Planck scale. Both model classes also include gravitons. Large Extra Dimension Models set the fundamental scale of nature to be the electroweak scale, and account for the weakness of gravity with its propagation into a Large Extra Dimension [32]. In Randall-Sundrum models with small extra dimensions, the disparity between the scales is due to an exponential warping of spacetime [33]. In both models, if gauge bosons are allowed to propagate into the extra dimensions, they can form Kaluza-Klein (KK) excitations. These KK excitations would appear as heavy versions of the existing gauge bosons, and would look like a Z' . A traditional Z' (originating from an additional $U(1)$) can also arise when compactifying the higher-dimensional metric.

String theories that include gravity also produce additional $U(1)$ groups. Superstring theories that rely on the $E_8 \times E_8$ of $SO(32)$ can be broken to low energy Supersymmetry. Depending on the compactification, the rank of the SUSY group is generally greater than five, resulting in at least one Z' [29].

Mentioned in passing in the preceding discussion is the fact that dark matter candidates often accompany a Z' . These candidates arise in anomaly cancellation, or can be a direct results of the symmetry breaking mechanism. This general feature of Z' models provides a phenomenological motivation for Z' searches. The weakly interacting dark matter of these models provides a promising avenue for collider-based dark matter searches, and offer both direct and indirect search channels for dark matter candidates.

Given the broad motivation for generic Z' models, the physics program of the Large Hadron Collider has an extensive selection of the Z' searches. Figures 2.4 and 2.5 show the current set of exclusion limits for various BSM searches and Z' final states in the Compact Muon Solenoid experiment.




Figure 2.4: Summary of the Compact Muon Solenoid Experiment’s “Beyond Two (2) Generations” (B2G) physics analysis group 2016, 2017, and 2018 Z’ results [34]. The B2G group is a BSM search group in CMS that specializes in intermediate vector bosons in the final state.




Figure 2.5: Summary of the Compact Muon Solenoid Experiment’s “Exotica” (EXO) physics analysis group 2016, 2017, and 2018 Z’ results [35]. The EXO group is a BSM search group in CMS that searches generally for exotic signatures.

With the ubiquity of a Z' in BSM models and their diverse decay mechanisms, evidence for a Z' could be hiding in unexplored signatures at the LHC.

2.2.2 Leptophobic cascade decay

A leptophobic Z' – forbidden to decay directly to SM leptons – is a candidate for such a hidden Z' . Leptophobic models can decay to two SM quarks, but dijet searches provide relatively weak limits on heavy Z' models due to the dominant and irreducible QCD background [36].

While none of the generic models in Section 2.2.1 specifically predict a leptophobic Z' , the leptophobic condition can be added as long as the model remains anomaly




Figure 2.6: Triangle diagrams of an axial current vertex.

free. Anomalies arise in chiral models with additional $U(1)$ groups due to effects of the triangle diagrams in Fig. 2.6. When these diagrams originate from an axial-current interaction, their vertex avoids renormalization. As a result, the axial current is not conserved. Naturally, this is not physical, and must be avoided. In 1972 Gross and Jackiw stated that the consistency condition of chiral gauge theories for a representation of R is given as

$$\mathcal{A}^{abc} = \text{tr}[t^a \{t^b, t^c\}] = 0, \quad (2.58)$$

where the $\mathcal{A}^{abc}$ is a trace over the group matrices in the representation [37]. For example, in the SM $SU(2)$, this condition becomes

$$\mathcal{A}^{abc} = \frac{1}{8} \text{tr}[\sigma^a \cdot \delta^{bc}] = 0, \quad (2.59)$$

where σ^a are the Pauli matrices, and δ^{bc} is the Kronecker Delta. The SM doublets also couple to the SM $U(1)$, giving

$$\mathcal{A}^{bc} = \text{tr}[Q \{ \frac{\sigma^b}{2}, \frac{\sigma^c}{2} \}] = \frac{1}{2} \text{tr}[Q] \delta^{bc}, \quad (2.60)$$

where Q is a matrix of the hypercharges. Given the hypercharges of the fundamental particles, the trace in Eq. 2.60 only vanishes if there are equal numbers of quarks and leptons: a fact of the SM. There is no a priori reason that the cancellation happens within a generation, but it does for the SM. Thus, the SM EW interaction is anomaly free.

An additional BSM $U(1)_{Z'}$ gauge group must satisfy the condition in Eq. 2.58. In the case of a leptophobic Z' , anomaly-free models arise with only the SM fermion content when there are generation dependent charges in the new symmetry group. If one only considered the sum of the quark charges in the new symmetry representation, each generation would have to have different charges, to have the trace analogous to Eq. 2.60 vanish. When one forsakes the generational independence, large flavor changing neutral currents (FCNCs) arise, which are unobserved in nature. Thus, to have a chiral-anomaly free model without FCNCs, new fermions,

field	spin	$SU(3)_C$	$SU(2)_W$	$U(1)_Y$	$U(1)_B$
$\bar{L}_L$	1/2	1	2	$-1/2$	-1
L_R					+2
$\bar{E}_L$	1/2	1	1	-1	+2
E_R					-1
N_L	1/2	1	1	0	+2
N_R					-1
ϕ	0	1	1	0	+3

Table 2.1: Fields that interact under the new $U(1)_B$ charge in the minimal anomalon model and their associated charges. $SU(3)_C$, $SU(2)_W$, and $U(1)_Y$ represent the SM groups, with the associated columns showing the strong, weak isospin, and weak hypercharges associated with each of the new particles.[38].

referred to as “anomalons”, must be included to cancel the chiral anomalies that arise with generation-independent charges. These new fermions together with the SM quark content satisfy the consistency condition in Eq. 2.58. Anomalons must be chiral with respect to the new gauge group, but vector-like in respect to the SM groups to respect the current LHC limits on additional generations of fermions [38].

The anomalons open an additional decay channel for the Z' : a channel potentially more sensitive than the QCD-dominated dijet region.

This thesis presents a search for a “baryonic” leptophobic Z' , denoted Z'_B [38]. A baryonic Z' is produced from a $U(1)_B$ gauge symmetry where the SM fermion charges are proportional to the baryon number, and are therefore generation independent. Minimally, the anomalons consist of three color-singlet SM-vectorlike fermions. These particles form a weak doublet $L = (L^\nu, L^e)$, and two weak singlets E and N . The E and the L^e are electrically charged, while the L^ν and the N are neutral. The anomalon content and quantum numbers are listed in Table 2.1. The dijet width of the Z'_B is given as

$$\Gamma(Z'_B \rightarrow t\bar{t}) \simeq \frac{1}{5}\Gamma(Z'_B \rightarrow jj) = \frac{g_z^2}{144\pi}M_{Z'} \quad (2.61)$$

where g_z is the coupling associated with $U(1)_B$, $M_{Z'}$ is the mass of the Z'_B (assumed to be $\gg m_t$, in this case), and j represents any hadronic jet. The width going to anomalons is given as

$$\Gamma(Z'_B \rightarrow N\bar{N}) = \frac{5g_z^2}{96\pi}M_{Z'}\left(1 - \frac{4M_N^2}{M_{Z'}^2}\right)^{1/2}, \quad (2.62)$$




Figure 2.7: Illustrative decays of the baryonic Z' to anomalons.

with similar expressions for the other anomalons. Depending on the mass-splittings between the Z' and the anomalons, the total anamol branching fraction can be as much as $5/6$. The Z' is given mass by the addition of a scalar field ϕ via the Higgs Mechanism, which contributes mass terms to the anomalons via Yukawa couplings. Each anamol also has Yukawa couplings to the SM Higgs. Due to the coupling to the Higgs field, the anomalons in table 2.1 are not the mass eigenstates. Minimal mixing of the anomalons that share the same electric charge preserves the EW symmetry. This gives the neutral left-handed mass eigenstates as

$$\begin{pmatrix} N_{S_L} \\ N_{D_L} \end{pmatrix} = \begin{pmatrix} \cos \theta_N & -\sin \theta_N \\ \sin \theta_N & \cos \theta_N \end{pmatrix} \begin{pmatrix} N_L \\ L_L^\nu \end{pmatrix}, \quad (2.63)$$

where θ_N is the mixing angle. The right-handed states are given by an orthogonal transformation. The charged anomalons have similar mixing, and produce analogous E_D and E_S states.

The anomalons decay via a cascade of lighter anomalons and SM bosons. In this minimal anamol model, the lightest anamol N_S is stable, ending the cascade, and offering a dark matter candidate. The heavier, neutral anamol decays to a neutral SM boson (a Z or a H , if kinematically allowed) and a stable N_S . The charged anomalons add one more tier to the cascade: the E_D decays to a W and an N_D . Figure 2.7 shows example Feynman diagrams for these decays.

The anamol decay channels of a baryonic Z' offer a yet unexplored phase space of Z' decay. The SM bosons in the final state and the kinematics of a cascade decay provide experimental signatures that new physics searches can exploit to enhance sensitivity over classic dijet searches that are dominated by QCD backgrounds. The instrumentation and analysis methods used to perform this search will be described in the following chapters.

Chapter 3

The CMS Experiment and the LHC

This chapter discusses the experimental facilities that make the work presented in this thesis possible. The Large Hadron Collider (LHC) at CERN provides the proton-proton collisions and the Compact Muon Solenoid (CMS) Experiment records the products of those collisions at one of the LHC collision points. The LHC is described in Section 3.1, with special emphasis on the design parameters in Section 3.1.1, the technology in Section 3.1.2, and finally the delivered beam in Section 3.1.3. The main components of CMS are described in Section 3.2. The coordinate system and trigger are described in Sections 3.2.1 and 3.2.7, and the subsystems are described in Sections 3.2.2 through 3.2.6, with the order indicating the subsystems' distance from the interaction point. The upgrade of one subsystem of CMS, the Hadron Calorimeter, is discussed in more detail in Chapter 4.

3.1 The Large Hadron Collider

The Large Hadron Collider was designed to discover the Higgs Boson and to open the frontiers beyond the Standard Model in the high energy regime. To do this, a 14 TeV center of mass proton-proton collider was designed. CERN, the European Organization for Nuclear Research, provided the perfect home for the LHC. The existing accelerator complex and tunnel at CERN crossing the French-Swiss Border served as the base and inspiration for many of the design elements of the LHC, as well as the main argument for the accelerator's presence at the laboratory. CERN was originally founded by twelve European countries after WWII to promote peace in Europe and to pool resources to allow for the construction of a world class experimental facility; the construction of the LHC transcended this goal. The LHC as we know it was made possible by the first instance of large, non-member state contributions to an accelerator facility, setting the example for a large scale international particle physics collaboration.

For the physics mission of the LHC, the colliding beams must not just be high in energy, but they also must have a high luminosity. The energy demands and the

tunnel restrictions point to a hadron collider, which can withstand the synchrotron radiation losses incurred at high energies. The LHC can collide heavy nuclei atoms as well as protons. The proton physics results are the focus of this thesis; therefore, the heavy ion operation will not be discussed. At LHC energies, the collisions are primarily between the sea-quarks in the proton nucleus, removing the benefit of a proton-antiproton machine to achieve large-scale quark-antiquark annihilation. The luminosity also demands a proton-proton machine. The design parameters pushed the extant superconducting acceleration and magnet technology to new cryogenic frontiers, making the LHC not just a feat for particle physics, but a feat for accelerator physics as well.

3.1.1 Beam Properties

In the search for new particles, two main beam properties dominate the reach of the collider: the center-of-mass energy, and the luminosity. Both parameters access rare physics in complementary ways. Increasing the center-of-mass energy provides more phase space for the production of heavy particles, while increasing the luminosity increases the number of interactions possible and thereby increases the rate of rare-process production. The LHC is designed to maximize both of these properties, while remaining within the design constraints of the existing services at CERN.

The center-of-mass energy, or the energy available when both of the beams collide, is straight forward: it is the sum of the individual beam energies. While the radio frequency (RF) cavities of the LHC provide the acceleration, the primary limitation on the collision energy achievable is the bending power of the dipole magnets that turn the protons through the LHC beam pipe. If the protons have too much energy, the magnets will be unable to bend them sufficiently to follow the beam pipe, and the protons will be lost and unavailable for collisions.

The luminosity of the machine [39] is best put in context of the rate of events (collisions),

$$N_{event} = L\sigma_{event} \quad (3.1)$$

where N_{events} is the number of events per unit time, σ_{events} is the cross section of the studied event, and L is the luminosity. The machine luminosity is dependent on many beam parameters, which are in turn determined by the design and operation of the machine. Nominally, the luminosity is given by

$$L = \frac{N_b^2 n_b f_{rev} \gamma_r}{4\pi \varepsilon_n \beta^*} F \quad (3.2)$$

where N_b is the number of particles per bunch, n_b is the number of bunches per beam, f_{rev} is the frequency of revolution, γ_r is the relativistic factor, ε_n is the normalized beam emittance, β^* is the beta function at the collision point, and F is the reduction factor that comes from the beam crossing angle at the collision. The emittance and the beta function are the two beam properties that require most of the artistry that goes into the LHC and injection chain design. Roughly speaking, the beam emittance characterizes the spatial distribution of the particles in phase space (a momentum space), and is invariant throughout the acceleration when normalized to beam energy, without the presence of the dissipative forces [40]. This invariance puts direct importance on the initial machine parameters, and the staging throughout the accelerator chain, as described in the next section. The beta function describes the transverse size of the beam. In the LHC, special magnets, described in Section 3.1.2, ensure a low-beta value at the high luminosity collision points.

3.1.2 Accelerator Chain and LHC Details

The LHC ring is the final stage of the CERN accelerator complex, with iterations of CERN's previous accelerators serving as the initial stages [41]. Figure 3.1 shows an illustration of the accelerator complex. Most of the beam parameters have to be carefully orchestrated prior to injection into the LHC ring. To prepare for running conditions expected by the LHC, the proceeding chain had to undergo substantial upgrades.

The source of the proton beam is a canister of H_2 that gets ionized to H^+ . The first accelerator in the chain is Linac 2, a linear accelerator. At the time of this thesis publication, Linac 4 is now the initial acceleration stage, but Linac 2 was the source for the data in this thesis. Originally, Linac 2 was injected with 750 keV protons from a Cockroft-Walton, but this was replaced with a Radio Frequency Quadrupole (RFQ2) to increase the intensity of the beam. At the end of Linac 2, the protons have an energy of 50 MeV, which makes them suitable for injection into the first circular stage in the complex, the Proton-Synchrotron Booster (PSB), where the electrons are also removed. The PSB determines the ultimate beam luminosity, and underwent an upgrade to deliver its current beam energy of 1.4 GeV to the next stage, the Proton Synchrotron (PS). This energy upgrade was needed to overcome space charge effects once the protons enter the PS. In the PS the 25 ns bunch spacing is implemented using the upgraded 20, 40 and 80 MHz radio frequency (RF) systems. In addition to forming the bunch train, the PS accelerates the protons to 25 GeV for injection into the Super Proton Synchrotron (SPS). The SPS is the penultimate acceleration phase, bringing the protons to 450 GeV before injection into the LHC. The SPS fills one LHC Ring, then the other, in cycles of roughly 3.2 s. Neglecting the time to switch between each ring, it takes about 9 minutes to fill the LHC.




Figure 3.1: The LHC accelerator chain and the experiments.[42]

The LHC is the final stage of acceleration and collides the beams [39]. The LHC has two counter-rotating proton beams that circulate in two separate rings. The 26.7 km LHC tunnel is buried between 45 and 170 m beneath the surface with an incline of 1.4%. The depth mainly allows for the land ownership of the surface to be independent of the tunnel beneath. The overburden also provides shielding for any extraneous interactions along the beam. The use of the tunnel originally built for the Large Electron Positron Collider (LEP) reduced the cost of the LHC and provided some of the more stringent design constraints. The tunnel has eight potential crossing points that are flanked by long, straight sections optimized for a lepton machine. Only four of the potential crossing points are used for collisions. The collision sites serve the ATLAS [43], CMS [44], LHCb [45] and ALICE [46] experiments. The other potential collision points serve as insertions points in the beam-line. The RF cavities for acceleration occupy one site, two sites are beam cleaning sites with collimators, and the remaining site houses the beam dump.

There are a total of 57 different varieties of magnet used in the LHC [41]. They can broadly be classified by use into three types:



Figure 3.2: Picture of cryodipole outside of the CERN Restaurant 1 patio.

1. bending magnets,
2. straight-section focusing magnets, and
3. insertion point magnets.

The iconic blue cryodipole (superconducting dipole and accompanying cryostat, Fig. 3.2) is the best example of the first type of magnet. The end of each dipole also features a collection of corrector magnets, to maintain the beam parameters. All in all, the LHC has 1232 superconducting NbTi dipole magnets, each nominally cooled to a temperature of 1.9 K to achieve the design field of 8.33 T. While the technology for the LHC dipoles is not new, this is the first use in combination with cryogenics pushed to this temperature. These magnets must be operated at or below 1.9 K for 8.33 T performance, making these magnets among the most complicated ones at the LHC. The bending regions of the tunnel have the smallest diameter of 3.7 m, which limits the size of the two-rings. Twin-bore dipole magnets are used, compressing the space needed for the two beam lines. This sacrifices flexibility because both beams are magnetically coupled.

The focusing quadrupoles in the straight sections of the LHC are made of the same materials, but operate at a higher temperature due to their reduced field requirements. Like the dipole magnets, the ends of the short-straight sections also have accompanying corrector magnets. Each quadrupole cryostat holds a set of quadrupoles that alternate focus in the x-y plane, called a lattice quadrupole. This keeps the beam in the correct envelope.

The insertion point magnets are a combination of many types of magnets, and this thesis will focus on the distinct features of the ATLAS and CMS insertion points, the high luminosity collision points. While the other sites feature their own interesting cadre of magnets, those parameters are unrelated to the physics focus of this thesis, and are therefore omitted. While superconducting magnets dominate most of the LHC, normal conducting magnets play a key role at the high luminosity collision points. Normal conducting magnets can withstand the higher radiation environments better than their superconducting counterparts. When the beams approach the experimental areas, pairs of two dipoles controlling the separation of the two beams bring the beams into and out of the collision orbits. For the first dipole upstream of the interaction point, a normal conducting dipole magnet is used. This first separation dipole is made of copper conductor and has a nominal field of 1.28 T. Right before the collision, a superconducting dipole is used. The low-beta triplet magnet resides between the two separation magnets. The low-beta triplet is a collection of four superconducting, single-bore quadrupoles and their associated corrector magnets. Without the focusing power of the low- β triplet, the luminosities sought at the collision points would not be possible. With a peak field of 7.7 T in the coil, the large bore to accommodate both beams, and the high heat and radiation loads experienced, the temperature requirements of the low- β triplet rival that of the main rings' dipole magnets.

To achieve the design collision energy of 14 TeV, superconducting RF cavities similar to that of LEP are needed. RF cavities are based on the same niobium sputtered onto copper technology. This composite material has several advantages over solid niobium, with the main improvement being the increased heat dissipation due to the presence of the copper. Unlike the dipole magnets, each of the LHC beams passes through its own RF cavity, owing to the larger tunnel and cavern at the RF insertion point. In total, the LHC has sixteen RF cavities (eight per beam), each powered by its own klystron.

3.1.3 Delivered Beam Characteristics

Run 1 of the LHC was the initial physics run, and took place between 2010 and 2012 with center of mass energies of 7 and 8 TeV. Run 2 of the LHC spanned 2015-2018, and the LHC delivered 163.54 fb^{-1} of data at 13 TeV. The Run 2 luminosity CMS recorded is shown in Fig. 3.3a. The design instantaneous luminosity goal of the LHC of $L = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ has already been exceeded twice over in Run 2, a remarkable feat of the LHC team. While the design energy of 14 TeV has not yet been achieved for physics operation, we await the prospects of 14 TeV running in Run 3 of the LHC or HL-LHC. The LHC has successfully completed two periods of data taking. This thesis presents a physics analysis using the data recorded from 2016-2018 in the CMS detector.




Figure 3.3: Total delivered and recorded integrated luminosity (3.3a) and average number of pile-up interactions (3.3b) measured at CMS during Run 2 of the LHC. Blank periods in the integrated luminosity plot show the periods of Year End Technical Stop (YETS) when the machine is down for maintenance.[47]

The benefits of the high luminosity come with challenges. The concept of luminosity inherently leads to the conclusion that several independent collisions can happen within the same bunch crossing. These collisions are termed *pile-up*, and are generally considered background within the experiments. Figure 3.3b shows the average pile-up recorded over the course of Run 2 in CMS. The pile-up is directly proportional to the instantaneous luminosity, and its identification and subsequent subtraction drives several of the design goals and reconstruction algorithms of the CMS experiment.

3.2 The CMS Experiment

The CMS detector [44] is one of the two large, general purpose detectors at the LHC. Concentric cylinders of specialized subsystems surround the interaction point and comprise the central “Barrel” region of CMS. To capture all of the collision products, CMS also includes disk-shaped endcap portions of each subsystem, and is completed with forward detectors farther along the beam pipe. The characterization of a collision needs the position, energy, and momentum information of the resulting particles.

As particles emerge from the interaction point they deposit energy in the detector material as they pass. Most stop within the detector volume. A low density silicon pixel and strip tracker lies closest to the interaction point. The tracker’s role is to measure the trajectories of particles while minimally changing the particles’ paths. With this information, the particles can be traced back to a specific vertex in the interaction region to offset the impact of pileup. When combined with a strong

magnetic field, like the almost 4T field of the CMS solenoid (see Section 3.2.5), the tracking information also provides an independent momentum measurement of charged particles from their bending radius in the magnetic field. A more detailed discussion of the CMS tracker is given in Section 3.2.2.

Once the initial trajectories have been determined, calorimeters measure the energies of electromagnetic and hadronic particles. When particles interact with matter, they “shower” or produce secondary particles, which in turn repeat the process until all of the energy of the initial particles has been spent. In CMS, the calorimeters convert this energy into light. Ideally a calorimeter captures the entire shower. Particles where the electromagnetic interaction dominates, like electrons and photons, will deposit most of their energy in the electromagnetic calorimeter. The electromagnetic calorimeter is the next subsystem particles encounter when radially exiting the interaction point. Electrons mainly lose their energy through the bremsstrahlung radiation of photons. High energy photons mainly lose their energy via pair production of electrons and positrons. This symmetry implies that both particles can be detected the same way, with similar governing parameters. In the case of electrons, the metric used in calorimeter design is the radiation length, X_0 ,

$$X_0(\text{g/cm}^2) = \frac{716 \text{ g cm}^{-2} A}{Z(Z+1) \ln 287/\sqrt{Z}}, \quad (3.3)$$

where Z and A are the atomic number and the atomic weight of the calorimeter active material, respectively [48]. The radiation length is the average distance the electron travels in the material until its energy is reduced to $1/e$ of its original value, where e is a chosen parameter. The ECAL precedes the hadron calorimeter (HCAL) to have precision electron and photon measurements. The radiation length of electrons is on average shorter than the analogous quantity for high energy hadronic interactions, the interaction length. The first layer of any HCAL would therefore primarily serve as an ECAL, so having specific instrumentation increases the electromagnetic resolution.

Nuclear interactions initiate hadronic showers, so the shower parameters are roughly defined by the characteristic length for hadrons to interact with the nucleons in the hadron calorimeter [49]. Hadron calorimeters also detect neutral hadrons, allowing for a complete characterization of the energy landscape. If the hadron calorimeter is hermetic, the presence of invisible particles, like neutrinos, can be inferred from the energy imbalance in the transverse plane of the detector. The momentum imbalance is referred to as *missing transverse momentum* ($\vec{p}_T^{\text{miss}}$) or *missing transverse energy* (MET). Sections 3.2.3 and 3.2.4 provide more detail on the CMS electromagnetic and hadronic calorimeters, respectively.

Muons are typically minimum ionizing particles and can escape the calorimeters due to their long lifetime and low energy deposition. The muons therefore pass through a

second layer of tracking in the muon systems. The CMS muon systems, described in detail in Section 3.2.6 are a combination of tracking and timing detectors. Like the silicon tracker, momentum measurements of the muon can be completed in this second stage and increase the resolution when combined the silicon tracker measurement. The timing is important to match the muons to the proper LHC bunch crossing, and to allow for a robust muon trigger.

3.2.1 Coordinate System

First, it is necessary to introduce the terminology to describe the structure of CMS. The CMS Experiment has adopted a right-handed coordinate system. The origin of this coordinate system is the nominal interaction point at the center of CMS. The z -axis aligns with the beam, with the positive direction towards the Jura mountains, or anticlockwise along the LHC ring. The x and y transverse plane is parallel to the endcap disks of CMS, with the positive x axis pointing towards the center of the LHC ring. Depending on the context, different coordinate systems are more convenient to describe CMS. Strictly for detector description, cylindrical coordinates are most convenient. The azimuthal angle ϕ rotates in the xy -plane from the positive x -axis. To describe particles, pseudo-spherical coordinate systems are more convenient, since the particles can emerge isotropically from the interaction point. Here the azimuthal angle remains the same. In classical spherical coordinates, the angle between the z -axis and the xy -plane would be given by θ , and in CMS this convention remains, but it is more convenient instead to use the pseudorapidity,

$$\eta = -\ln(\tan(\theta/2)). \quad (3.4)$$

The relationship between pseudorapidity and θ is shown in Fig. 3.4. In reality, the “interaction point” is a roughly 1 m long region in the center of CMS, and particles collide all along this region. Each interaction is referred to as vertex, reconstructed by the paths left in the tracker, described in Section 3.2.2. Once this location along z is determined, the kinematic description of particles are then referenced to this location as the origin. Since the particles are moving at relativistic speeds, to determine the angle between them, but in the yz -plane, a Lorentz invariant quantity is needed. The classical θ no longer suffices, but the *rapidity* offers this feature. When the energy of the particle is much greater than the mass, the pseudorapidity is a good approximation of the rapidity.




Figure 3.4: A diagram of the relationship between η and θ . Pseudorapidity values mirrored across the z -axis are negative.[50]

3.2.2 Tracker

Closest to the collision point in CMS lies the silicon tracker. The tracker's role is to capture the paths of the produced charged particles, trace them back to their respective vertices, and measure their bending radius, courtesy of the 4T magnetic field that uniformly bathes the tracker volume. Silicon solid-state sensors offer the speed and radiation tolerance demanded of the CMS tracking system. The devices must be fast enough to inform the high level trigger of CMS. In such close proximity to the collision point, the tracker also accrues radiation damage over time and cryogenic operation is the primary tool to mitigate these effects. The cryogenics are the final important aspect of the tracker design. The cryogenic services and tracker support structure must have an exceedingly low material budget to minimize energy deposition before the calorimeters.

The CMS tracker uses both silicon strip and silicon pixel sensors. While strips are more cost effective, and have simpler readout needs, the high particle flux in the immediate vicinity of the beam pipe requires the finer resolution of a silicon pixel detector. Figure 3.5 shows a schematic of the pixel and strip locations in the CMS tracker. The silicon pixel detector present for the 2016 data-taking period was made of three barrel layers 53 cm long and positioned at radii of 4.4, 7.3, and 10.2 cm. The central region was called “BPix”, and was accompanied by 4 endcap disks, two on each side, termed “FPix”. These disks were placed along the beam axis 34.5 cm and 46.5 cm away from the interaction point. Their radii were optimized to extend the pixel detector's coverage to $|\eta| < 2.5$, with at least three hit layers possible within this

region. All together, the original CMS pixel detector had 66 million pixels, grouped in 1440 pixel modules. Between the 2016 and 2017 data-taking periods, the Phase 1 upgrade of the CMS Pixel detector was completed. The Phase 1 pixel incorporates the same sensor technologies as the original design, but adds one additional barrel layer, additional disks to each endcap, and exchanges the support structure and cooling for a newly available carbon fiber support and carbon dioxide cooling system [51]. The new barrel layers lie at radii of 2.9, 6.8, 10.9, and 16.0 cm. The total silicon area increased from 1.1 m^2 to 1.9 m^2 , but new mechanics keep the material budget the same. Figure 3.6 shows a comparison between the original (Phase 0) and the upgraded (Phase 1) pixel detector.

The Phase 1 pixel upgrade kept much of the same sensor design as the original. Each pixel is an n+ pixel on a n-substrate, $100 \times 150 \mu\text{m}^2$ in area. Within the pixel barrel, the electron drift is perpendicular to the magnetic field. The Lorentz force on the drifting electrons causes the electrons to produce signal in adjacent pixels, a phenomenon known as charge spreading. This increases the resolution of the device, since the analog readout allows for the peaking pixel to be clearly visible amongst a large, and therefore easily identifiable, cluster. To get the same benefits in the forward disks, the pixel faces are rotated 20° away from the normal, in a turbine structure. This concept was taken a step further in the upgraded design, with the innermost FPix rings tilted 12° with respect to the beam line in an inverted-cone structure. On average, incoming particles will hit the forward pixels normally, maximizing the drift distance for the produced electrons, while also having non-zero Lorentz force.

The strip detectors are placed radially beyond the pixel volume. The strips are further classified into two parts: the inner tracker (barrel and disks, TIB and TID, respectively), and the outer tracker (barrel and endcap, TOB and TEC, respectively). The inner tracker is made of micro-strip sensors $320 \mu\text{m}$ in thickness and with a cell size of $10 \text{ cm} \times 80 \mu\text{m}$. To reduce the number of read-out channels in the outer tracker, the strips in the 55 cm to 110 cm radial region can be as large as $25 \text{ cm} \times 180 \mu\text{m}$, with a $500 \mu\text{m}$ thickness. The strips are *p-on-n* type semiconductors, and the choice of the larger thickness in the outer tracker region was to reduce the capacitive noise that scales with the increase in the strip length.

In addition to the combinatorial and geometric complexity of the tracker, all of it must be kept around -10°C . With such low temperatures, the humidity of the tracker must be kept very low. This requirement affects the rest of the subsystems housed in the magnet volume. Each subsystem following the tracker must maintain the tracker's vapor barrier, dictating the installation and operation of the rest of the barrel region of CMS.




Figure 3.5: A cross-section schematic of the Phase 0 CMS tracker.[44]




Figure 3.6: A cross-section comparison of the Phase 0 CMS pixel tracker (turquoise, bottom) and Phase 1 pixel tracker (blue, top) .[51]

3.2.3 ECAL

Once particles exit the tracker and pass through the tracker's thermal screen, they enter the CMS Electromagnetic Calorimeter (ECAL), a homogeneous and hermetic

lead tungstate (PbWO_4) crystal calorimeter [44]. The ECAL is responsible for the measurement of the energies of particles that interact electromagnetically, with an emphasis on electrons and photons. Since photons do not bend in the magnetic field, the ECAL is responsible for photon energy measurement. The ECAL barrel is made of 61,200 crystals, each with a flared shape, with a smaller face in the direction of the interaction point, and a larger face in the rear, to maintain hermeticity. The shapes slightly vary with η , but on average have an interaction-point-facing size of $22 \times 22 \text{ mm}^2$, and a rear-facing size of $26 \times 26 \text{ mm}^2$. Each crystal encompasses 25.8 radiation lengths at 230 mm long. Each crystal axis is tilted 3° from the interaction point normal to ensure that the spaces between crystals do not align with particle trajectories. The endcaps extend the calorimeter coverage to $|\eta| < 3$. The ECAL endcaps each have 7324 crystals. Unlike the barrel, all of the crystals are uniform with a length of 220 mm ($24.7 X_0$), a front-face area of $28.62 \times 28.62 \text{ mm}^2$ and a rear-face area of $30 \times 30 \text{ mm}^2$. Similar to the barrel, the crystal faces are not focused at the interaction point, but are focused at a point 1.3 m beyond the interaction point.

Lead tungstate crystals offer the granularity, speed, density, and radiation hardness required of the CMS ECAL. First, the short radiation length (0.89 cm) and high density (8.28 g/cm^3) allow for a compact calorimeter to fit in combination with a hadron calorimeter in the volume of the CMS solenoid. The Molière radius of 2.2 cm offers high-granularity, reflected in the crystal geometry. Finally, the bulk of the scintillation light is produced within 25 ns, aligning with the LHC bunch crossing timing demands. Unfortunately, the crystals are very responsive to temperature. Therefore, the crystals and their corresponding light sensors are cooled to 18° C within a tolerance of 0.05° C .

Both the ECAL barrel and endcap have custom-designed photodetectors. The primary requirement is gain, or signal amplification, to accommodate the moderate levels of scintillator light. The ECAL barrel is fitted with solid-state avalanche photodiodes (APDs), operated at a gain of around 50. Two APDs serve each crystal. The ECAL endcaps use vacuum phototriodes (VPTs) to withstand the higher radiation environment. The VPTs were designed for operation in the CMS magnetic field. Even though they are less sensitive per unit area than the barrel APDs, the vacuum tubes accommodate by covering a larger surface area on the back of each endcap crystal. On average, signals are also larger in the endcaps. The VPTs are operated at a gain of around 50, but require significantly higher bias voltages.

3.2.4 HCAL

Following the ECAL is the CMS Hadronic Calorimeter (HCAL) [44]. The HCAL is a sampling calorimeter comprising four distinct subdetectors: the HCAL Barrel (HB),




Figure 3.7: Labeled locations of the the HCAL subdetectors.

the HCAL Endcaps (HE), the HCAL Forward (HF), and the HCAL Outer (HO). The locations of each of the HCAL subdetectors are shown in Fig. 3.7. The CMS HCAL is the only subsystem on CMS capable of measuring neutral hadrons, and therefore must be hermetic, in order to reconstruct missing transverse momentum. The HCAL barrel and endcap are made of alternating layers of brass and scintillating plastic, and together cover up to $|\eta| < 3$. The HCAL barrel is constrained to fit radially between the ECAL ($R = 1.77$ m) and the inner magnet radius ($R = 2.95$) which only allows 5.82 interaction lengths of absorber at the very center. To capture any remaining hadronic shower, the HO sits outside of the magnet volume, using the magnet as absorber. The HCAL Forward extends the HCAL's coverage to $|\eta| < 5$. The HF experiences the highest radiation doses within HCAL, and therefore uses radiation tolerant quartz fibers as the active material.

The HCAL barrel's absorber, except for the first and last layers, which are made of stainless steel for support, are made of large brass plates running parallel to the beam line. Brass was chosen for several reasons [52]. First, the long absorber plates must be non-magnetic, to avoid deformation under the 4T solenoid field. Second, the material must be sturdy enough to support its own weight and provide the structural support for the rest of the HB active material. Third, the material must have a short interaction length to make the most out of the compressed space between the magnet and the ECAL. Copper, and later brass, was settled on to meet the further constraints of affordability and ease of handling. Two types of scintillating plastic are used in the HCAL Barrel. For the layer directly behind the ECAL, and before the first steel absorber plate, 9 mm thick Bicron BC408 is used. This layer is responsible

for measuring the hadronic showers that originate in the ECAL, using the ECAL as an absorber. The rest of the layers are 3.7 mm thick Kuraray SCSN81. This combination of materials was chosen to enhance the performance of a depth-segmented readout [52], with the more luminous BC408 as a stand-alone layer. While this configuration was absent for Run 1 and Run 2 of the LHC, the HCAL Phase 1 upgrade utilizing this feature was completed in 2019, and is detailed in Chapter 4. The HCAL Endcap shares the same absorber and scintillating material configuration, but differs in mechanics and implementation, due its placement perpendicular to the beam pipe, and since it must be mounted to the endcap return yoke of the magnet.

The HCAL Outer (HO) uses additional 10 mm layers of the BC408 scintillator as a tail-catcher outside of the solenoid to fully capture hadronic showers. The HO is mounted as the first layer of each of the five magnet return yoke rings, which are described in more detail in section 3.2.5. At $\eta = 0$ HB has the fewest interaction lengths, so the corresponding HO ring is augmented by 2 layers of scintillator, and an additional 19.5 cm thick piece of iron absorber. With the addition of HO, the HCAL depth at a minimum is extended to 11.8 interaction lengths in the barrel.

The final subdetector of HCAL is HF, which covers $3 < |\eta| < 5$. It is distinct from the other HCAL detectors in that it does not use scintillating plastic as the active material, but instead uses quartz fibers as a source of Cherenkov light. This design choice reflects that HF receives almost 8 times the particle flux of the other subdetectors. Steel absorber is used in HF, and the fibers are run in grooves within the absorber itself. HF has two longitudinal sections, one with fibers spanning the entire detector, and one with the fibers starting 22 cm from the front of the detector. This is used to distinguish the electromagnetic and hadronic portions of the shower. The hadronic showers should span the whole detector, while the electromagnetic should focus in the first 22 cm.

3.2.5 Magnet and Flux Return Yolk

The CMS solenoid first impresses with its scale. The 6.3 m completely-open internal diameter and 12.5 m length make the CMS solenoid the largest 4T magnet ever built. The “cold mass” of the magnet, the components held at cryogenic temperatures and producing the field, total 220 t. While the dimensions alone are unique, the technical demands of such parameters truly distinguish the CMS solenoid. There are two primary design achievements of the solenoid. First, to achieve the 4T field in the inner volume, the CMS magnet consists of 4 times the winding layers of most previous HEP collider experiment magnets, for a total of 4 layers. Second, to keep a slim radial extent and to accommodate the large stored-energy to mass ratio, the conductor itself had to be innovative. Traditionally, the conductor has exclusively external support, but the NbTi conducting cable used in the CMS magnet has an Aluminum core,



Figure 3.8: A photo of a yoke endcap disk, the central wheel with the vacuum tank, and an additional barrel wheel during magnet assembly in SX5.[44]

providing the structural support within the current-carrying cable itself. This feature allows for the free-space in the internal diameter of the magnet, and for its slim profile.

While the magnet is CMS's namesake, the return-yoke for the magnetic field gives CMS its signature look, and modular design. The steel return yolk, shown in Fig. 3.8, painted red, consists of 5 barrel wheels and 6 endcap disks, totaling 10,000 t. The solenoid and its cryostat are mounted on the central wheel. Each section is mobile through the help of large air and grease pads, allowing for relatively straightforward access to the embedded subdetectors, with the wheels able to move outward. The solenoid vacuum tank and cryostat thread the inner diameter of wheels. The largest distance the segments can be offset is about 11 m, and at full speed, this can be done in about 1 hr. At final closure, the yoke can be aligned within 2 mm of its ideal position. The endcap calorimeters and the muon system cathode-strip chambers (CSCs) are hung from the endcap disks, and yoke wheels provide the structural support for HO and the barrel muon systems.

3.2.6 Muon System

The muon systems of CMS are the last layers of detector traversed when exiting the interaction point. Like the rest of CMS, the muon systems are divided into barrel and endcap regions. Since muons also enter CMS's name, it is worth noting why. Muons offer one of the cleanest windows into the Higgs discovery channel, with the

Higgs to four muon decay standing out clearly amongst the background. With the calorimeters stopping most of the hadron activity, the muon systems are free to select and reconstruct the clean trajectories of the muons that survive.

The muon systems provide three main roles in CMS. First, they provide muon identification, capturing them as they leave the calorimeters. Second, they provide another measurement of the muon momentum, which when combined with the tracker measurement gives a higher resolution. For the Phase 0 system, a high momentum muon with an energy of 1 TeV had a momentum resolution of around 5% when combining both tracker and muon system information. Finally, to be sensitive to rare decays going to muons in the final state, the muon system must also take part in the trigger, demanding a fast readout. Living on the outermost regions of CMS, the muon systems must cover a large area (22,000 m²) which incurs the further demands that the chambers must be inexpensive and robust. To tackle this, the muon system of CMS has three different types of gaseous detectors: drift tubes (DTs) for the barrel, cathode-strip chambers (CSCs) for the endcaps, and resistive-plate chambers (RPCs) for the trigger.

The drift tubes are nestled in the gaps of the barrel wheels of the flux return yoke, which imposes strict geometric requirements. The DTs are organized in four radially-displaced stations, with the first three stations equipped to measure $r - \phi - z$ information, and with the final station only recording $r - \phi$. Together, the DT covers $|\eta|$ out to 1.2. Within each station, the chambers are offset by a half-cell to ensure complete coverage. Each drift tube has a cross section 13×42 mm², and a length determined by the orientation of the tube: parallel to the beam line for ϕ position information, or perpendicular to the beam line for z -coordinate measurement. The tubes are filled with a gaseous mixture of 85% Ar and 15% CO₂, which acts as the gain medium. In addition to the anode wire in the interior of the tube, and the cathode along the sides, additional electrodes run along the top and bottom of the tube, as shown in Fig. 3.9 (left), to keep the electric field uniform. This uniformity is important to maintain a linear relationship between the distance from the wire and the drift-time of the electrons to allow for a time resolution of a few nanoseconds for triggering purposes.

A different technology was chosen for the endcaps: the cathode-strip chambers (CSCs). Each CSC has a petal-like structure, mounted on the endcap disks of the magnetic field return. The CSCs extend the muon system triggering and tracking capability to $|\eta| < 2.4$, with a stub finding design efficiency of 99% available for the L1 trigger, and an offline $r - \phi$ resolution between 75 and 150 μ m, depending on the chamber location. The CSCs are sandwiches of anode wire planes and cathode strip panels to make a total of six active gaps filled with a 40% Ar + 50% CO₂ + 10% CF₄ gas mixture. The anode wires run azimuthally, and the strips run perpendicular to the wires, to give $r - \phi$ position. Figure 3.9 (right) shows a simplified diagram of the CSC design. For the majority of the CSC chambers, 50 μ m diameter gold-plated tungsten





Figure 3.9: Left: Stylized diagram of a functioning cross-section of a drift tube. Right: Stylized diagram of a CSC showing the azimuthal wires, and the perpendicular strips.[44]

wires are wound around the anode planes, the longest of which is 3.4 m. The 80 strips are milled into the cathode planes, with the strips thinnest at the narrow end of the chamber, and widening as the chamber widens. While most of the CSC chambers exist with the wires completely perpendicular to the beam line, the section closest to the interaction point lies within the magnetic volume, and therefore experiences the strong magnetic field uniformly in the z direction (see Fig. 3.10). In these chambers, the wires are tilted at an angle of 29° , to have the electron drift parallel to the strips, to compensate for the Lorentz drift.

Since the muon system's role in the trigger is paramount, a third muon detector is included in CMS: the resistive plate chambers (RPCs). The RPCs unambiguously can match muon hits to the correct bunch crossing within the 25 ns window, and provide rough positional information that can be augmented by the finer granularity detectors. The large parallel plate design creates a uniform electric field, allowing for a better timing resolution than the other two flavors of muon detectors, which have less uniform fields. An RPC in CMS has two, 2 mm wide gaps that share a common set of read-out strips in between. The resistive plates are made of 2 mm thick Bakelite and vary in dimension depending on the installed location. The RPCs flank the drift tubes in the barrel, as shown in Fig. 3.11. The endcap RPCs are mounted directly on the endcap disks, in three stations. The first two cover both sides of the disk closest to the interaction point, and under the CSCs for the side facing the interaction point. The last station covers the interaction-point-facing side



of the second disk. All together, the endcap RPCs extend the coverage to $|\eta| < 1.6$. Like the rest of the muon systems, the RPCs use a gas as the gain medium. The gas mixture is made of 96.2% $\text{C}_2\text{H}_2\text{F}_4$, 3.5% $i\text{C}_4\text{H}_{10}$, and 0.3% SF_6 . Water vapor is also added to keep the humidity at 45% to preserve the Bakelite resistivity.

3.2.7 Trigger System

The CMS trigger system lies on the interface between the detector design, and the offline reconstruction. At the nominal proton-proton physics run design, CMS will receive collisions every 25 ns, or at a frequency of 40 MHz. This rate, increased further by pileup, is impossible to maintain through the entire data acquisition chain. The physics goals of CMS however are low-rate phenomenon, so the solution to the data inundation is clear: save less of the data. To do this, CMS uses a trigger system to only write data that has been deemed interesting, based on a few rudimentary tests. The CMS trigger system is two level. This first level, termed Level 1 (L1), is a hardware-implemented trigger, nominally designed to reduce the data rate from 40 MHz to 100 kHz [44]. The Level 1 trigger acts on rudimentary energy deposits from the calorimeters and basic tracks from the muon system. The second level is the bridge to offline reconstruction, the High Level Trigger. The High Level Trigger algorithms have access to the entire high resolution readout from all of the detectors, and can make more sophisticated decisions. For this combination, the front end electronics must buffer the data from each collision for $3.5 \mu\text{s}$, until the L1 decision is made, and the data is delivered to the High Level Trigger.

The L1 decision is made by aggregating information from the calorimeter and muon systems. The architecture of the L1 system is shown in Fig. 3.12. Each subsystem is responsible for generating trigger primitives, which collect either the track information in the muon systems or the calorimeter energy deposits, and the assigned bunch crossing. Within the subsystem group, a regional trigger takes that information for a spatial region of the detector, and makes an initial object identification. That regional trigger information is then ranked either at the calorimeter level or entire muon system level, and the regional trigger objects deemed most interesting are passed to the final step, the Global Trigger. At this point, the Global trigger makes the decision to accept or reject the event. For example, this decision algorithm at its most basic implements simple momentum cuts on the objects received. This decision, referred to as the L1A (level one accept), is then transmitted back to the readout electronics buffering the high resolution data, and the passing events are passed to the High Level Trigger.

The High Level Trigger has more flexibility than the L1 trigger because it has a longer time to make the decision, and then can use the full resolution with an approximation of the offline reconstruction. After the HLT decision, the overall data rate should be roughly 100 Hz [53]. The HLT operates efficiently because the reconstruction is only applied to the region of the detector that the L1 trigger has indicated. The HLT algorithms are periodically updated to align closely with the offline reconstruction. As new tools are implemented in the offline reconstruction, they can eventually be incorporated into the HLT.




Figure 3.12: Diagram of flow of L1 trigger.[44].

Chapter 4

The Hadron Calorimeter Phase 1 Upgrade

The Hadron Calorimeter plays an important role in the reconstruction of events in CMS, most notably as the only subdetector capable of measuring neutral hadrons. The hermetic design of the calorimeter allows for the reconstruction of missing transverse momentum, a feature that augments HCAL’s existing roles in the calorimeter-based L1-triggers and in electron identification. With the unprecedented luminosity conditions of the HL-LHC on the horizon and the shortcomings of the original front end electronics degrading performance, CMS decided to replace the HCAL front end during the “Phase 1 Upgrade.”

There are three phases of CMS, each preparing CMS for the next stages of LHC physics. Phase 0 is the originally installed system. Phase 1 is the system after upgrades installed during Long Shutdown 2 (LS2) of the LHC, between Run 2 and Run 3 of 13 TeV (13.6) TeV running. After the completion of Run 3, the LHC itself will transition to the High Luminosity LHC (HL-LHC). The running conditions of the HL-LHC dramatically differ from those informing the Phase 0 design. Phase 2 upgrades will complete CMS’s preparation for the high luminosity Run 4, and will be installed during Long Shutdown 3 (LS3).

My work has encompassed almost the entirety of the CMS HCAL Barrel upgrade: the first PCB tests, installation, operations, and finally solving a mystery that ultimately impacted a large part of the High Energy Physics community. The following sections describe my contributions to the upgrade effort. First, Section [4.1](#) introduces the HCAL design and upgrade goals. Section [4.2](#) describes my roles within the HCAL Barrel Upgrade. Finally, Section [4.3](#) recounts the VTRx (Versatile Link Transmitter/Receiver) failure investigation.

4.1 HCAL Endcap and Barrel Overview, and Upgrade Motivation

While all four of the HCAL sub-sections mentioned in Chapter 3 have been upgraded, it is convenient to group the HCAL Endcap (HE) and HCAL Barrel (HB) upgrades together; both detectors face the same challenges and feature similar solutions. Both HB and HE are sampling calorimeters with alternating layers of scintillating plastic and brass, and both have their data acquisition and control electronics situated on-detector, within the body of CMS. The electronics comprising the front end must meet several challenges: the electronics must meet the speed, bandwidth, and dynamic-range demands of LHC data-taking and operate within the confines of the CMS detector with the associated radiation, magnetic field, and spatial constraints.

4.1.1 HBHE Segmentation and Nomenclature

The HCAL Barrel [44] is divided into two halves. From an aerial perspective, HB plus (HBP) faces the Jura mountains and HB minus (HBM) faces Lake Léman. Each half consists of 18 identical wedges, with each wedge covering twenty degrees in ϕ . The HCAL Endcaps follow the same pattern: each endcap is subdivided into 18 matching wedges, the positive z-axis endcap is called HE plus (HEP) and the negative z-axis endcap is denoted HE minus (HEM). Each wedge corresponds to one Read Out Box (RBx) which holds the instrumentation for that wedge. The numbering and geometry is shown in Fig. 4.1. Numbering begins with 1, aligned the positive x-axis, and continues clockwise until 18 if facing towards the interaction point.

In the barrel, there are a total of 16 layers of absorber and 17 layers of scintillator, with the thicknesses varying longitudinally to balance performance and cost. Finer granularity is desired closer to the beam-pipe to capture shower development and to enhance particle identification, but the benefits of fine-segmentation wane deeper into calorimeter. The endcap has 17 layers of scintillator as well. Azimuthally, each wedge is divided into 4 parts. As one moves longitudinally through the HCAL Barrel, the layers are staggered to eliminate dead areas, as shown is Fig. 4.2.

Each wedge is further divided into “towers”. Each tower is a division that extrapolates the $(\Delta\eta, \Delta\phi)$ region of a readout tower in ECAL into the HCAL, providing the “CaloTowers” used for the Level 1 calorimeter trigger [44]. Each half-barrel has 16 η divisions, and each endcap has 13 η divisions (with a 14th “guard ring”). Tower numbering begins along the y-axis with tower 1, and increases with increasing η to tower 28 (29 is the guard ring). Tower 16 spans both HB and HE. Figure 4.3 shows the HCAL barrel and endcap segmentation and Phase 0 readout. In the Phase 0 system, towers 1-14 had no longitudinal readout segmentation, resulting in the loss of




Figure 4.1: The numbering begins on the inside edge of the LHC ring. For example, looking at HBM while standing on a cherry picker in front of HEM, HBM1 would be at 9 o'clock, and HBM10 would be a 3 o'clock. [44]




Figure 4.2: A diagram of an HCAL Barrel wedge to show staggered layers [44].




Figure 4.3: An yz -plane slice of the Phase 0 HCAL Barrel and Endcap. Towers and scintillating layers are indicated. Colored regions show the summed readout per-tower in the Phase 0 system. “FEE” indicates the location of the front end electronics, or the “RBx”.

any depth-dependent information. In the barrel, towers 15 and 16 had two readout layers. In the endcaps, the Phase 0 configuration had 2 readout layers in towers 18-27, 3 layers in 28-29, and no segmentation for towers 16-17.

The scintillating tiles follow the same segmentation. In the HCAL barrel, there are roughly 70,000 scintillating tiles. The tiles in a given longitudinal layer in a wedge are housed in a tray that slides between the absorber plates, allowing for removal in extreme circumstances. The collection of tiles in a tray is called a *megatile*. Each tile is routed for wavelength-shifting fibers to gather the light from the tile and carry it to the edge, where it is spliced to clear fibers. These fibers are then routed to the end of the tray, where a megatile cable (an optical fiber bunch) connects the fibers to the front end readout system. Figure 4.4 reveals the layered brass-scintillating panel structure of the HCAL barrel before installation.

Each five degree division in ϕ corresponds to one “read out module” (RM). Each RM is a card pack of the charge integration boards and light detectors responsible for the data acquisition. With the RM the light is combined into readout towers, measured as an analog pulse shape, converted to digital data, and piped off detector via an optical link. A RBx holds four RMs, a calibration unit, and control electronics.




Figure 4.4: An image of HBP prior to insertion into the magnet's vacuum tank above ground in SX5 in March, 2006. [54]. Image shows not only the scale of one half of the HCAL Barrel, but the layered structure of the absorber. This is no longer visible, as it has been covered by the cable trays of the tracker.

4.1.2 Phase 0 HCAL Challenges

In addition to the limitations inherent in the coarse read-out, the Phase 0 HCAL was subject to a plethora of other issues associated with the front end electronics [55, 56]. A key component of the Phase 0 system was the use of hybrid photodiodes (HPDs) as the optical transducers. The HPDs required a large bias voltage to operate, and were subject to electrical discharges that were augmented by the strong magnetic field of the CMS solenoid [55]. This effect, accompanied by unchecked photocathode migration [55], degraded the HPD performance and threatened the longevity of the transducers. In addition to the high-amplitude HPD noise, the Phase 0 HCAL also suffered from “RBx noise”, an unidentified noise signature that would affect an entire RBx, and resulted in an over-firing of the MET triggers [56]. While this was easily filtered, it hinted at unsolvable issues intrinsic to the Phase 0 electronics.

4.1.3 HL-LHC Demands

The future provides exciting challenges for CMS. With the approval of the High Luminosity LHC (HL-LHC) in 2016, the LHC community looks forward to energy frontier physics beyond the end of Run 3. The HL-LHC is designed to nominally multiply the initial LHC data-taking expectations by 10 times, and to increase the instantaneous luminosity by a factor of five [57]. These improvements are well within the realm of possibility, since the LHC has a history of achieving and exceeding its luminosity expectations. With the designed instantaneous luminosity of $1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ doubled during Run 2 of the LHC, it is natural to anticipate a fully-realized HL-LHC scenario. The CMS HCAL was designed for the nominal LHC luminosity, which translated to roughly 25 interactions per bunch crossing. Within the HL-LHC design margins, a maximum of 200 pile-up interactions per bunch crossing is expected[57]. Looking to this future shows that the CMS HCAL will have to adapt.

4.1.4 Phase 1 Upgrade Overview

The challenges of the HL-LHC can be met by the original design of the HCAL Barrel’s active material. Intrinsic to the sampling calorimeter’s design is a longitudinally segmented structure, allowing for up to seventeen distinguishable layers of readout. Figure 4.5 shows the Phase 1 Upgrade readout, with four readout layers in the barrel. In the high-pile-up environment of the HL-LHC, most of the background hadronic activity will be localized in the first layers, so increasing the readout segmentation gives an additional handle to suppress this background.

Naturally, replacing the existing readout electronics to accommodate the desired depth segmentation also replaces the noisy, extant front end, improving the prospects




Figure 4.5: An yz -plane slice of the Phase 1 HCAL Barrel and Endcap. Towers and scintillating layers are indicated. Colored regions show the summed readout per-tower of the upgraded Phase 1 system. “FEE” indicates the location of the front end electronics, or the “RBx”.

for the remaining LHC physics program. In the time between the Phase 0 and Phase 1 designs, Silicon Photomultipliers (SiPMs) became widely available and met the stringent demands of an on-detector phototransducer, offering a prime replacement for the problematic hybrid photodiodes in HB and HE. Magnetic-field tolerant, high gain, compact, low-bias, and moderately radiation tolerant, SiPMs should thrive in the years to come. The new front end electronics also feature new versions of the charge-integration and encoding (QIE) ASICs (application specific integrated circuits). These new chips are responsible for the charge collection and digitization of the SiPM signals. Additionally, the chips encode timing information, a new feature of the upgraded system that allows a new generation of timing-based triggers that can be optimized for exotic signatures.

The installation of the new front end electronics can only occur during long periods of shutdown, since CMS has to be opened to gain access to the RBxes. This leaves only the Year-End Technical Stops (YETS) and Long Shutdowns 2 and 3 (LS2 and LS3) to implement any changes. With the bulk of CMS subdetectors to undergo upgrades to face the HL-LHC during LS3, and since the Phase 0 system left much to be desired, the HCAL upgrades were installed on a rolling basis throughout the YETS in Run 2 and in LS2. This choice steadily improved HCAL’s performance throughout Run 2 with substantial upgrades to the HCAL Forward and HCAL Endcaps completed in 2016 and 2017, respectively. The HCAL Barrel upgrade was completed during LS2,

due to the access requirements of the central barrel. The Phase 1 upgrade to the Barrel will carry the CMS HCAL to the end of CMS's lifetime.

4.2 Lock, Stock, and HCAL Barrel

I took part in the final installment of the HCAL Upgrade, the HCAL Barrel. My HCAL journey began as a team member in the qualification of the QIE 11 boards for the charge integration and encoding of the new barrel front end at Fermilab. I followed the electronics to CERN, where I led the assembly and testing of the control electronics, and was an integral team member for the installation of the complete new front end electronics suite underground in the CMS cavern. After we connected the last megatile fiber, I transitioned into an avid operations team member, commissioning the new detector, and ensuring that an operational HB arrives at the start of Run 3, despite the setbacks due to the COVID-19 pandemic.

4.2.1 Phase 1 Barrel Front End

The basic structure of the Phase 1 front end and the Phase 0 front end are the same:

1. 4 read out modules (RMs)
2. 1 calibration unit (CU)
3. Clock, Control, and Monitoring Modules (CCMs)
4. Backplane to connect everything

but designing the upgrade in the mold of the existing system is a challenge. More channels must fit into the geometric footprint of the Phase 0 system, constrained by the aluminum shell of the RBx. In the case of the read out modules, this manifests in RMs that are larger than the aluminum RBx shell, protruding into the space between the shell and the exterior of the magnet's vacuum tank.

Furthermore, the design must accommodate the change in accessibility between the construction of the Phase 0 and Phase 1 systems. While Phase 0 was installed prior to tracker, the Phase 1 system now has the tracker cable trays bisecting RBx access. This poses a challenge for the removal of the Phase 0 front end, and the subsequent installation of the Phase 1. In Phase 0, one long backplane spanned the RBx, and with the cable tray in place, it was unclear if the new backplane of a similar length could be fit into the RBx around the tracker's cable trays. To address this, the Phase 1 barrel features a two-piece backplane. This is an excellent opportunity to increase




Figure 4.6: A schematic showing the naming, constituents, and organization of an HB RBx. Since HB RBxs are in 360°C, this specific orientation would only be found in the 6 o’clock position of either HB half.

the control granularity, so in addition to the two-half backplane design, two Next Generation Clock, Control, and Monitoring (ngCCM) modules were incorporated into the design, one per half-backplane. Figure 4.6 shows a schematic of the Phase 1 HB RBx.

4.2.2 Next Generation Clock Control and Monitoring

The Next Generation Clock, Control, and Monitoring Module (ngCCM) orchestrates the communication for the upgraded Phase 1 front end. For the HCAL Barrel, seventy-two HB ngCCMs are needed for operation. New control modules were needed to accommodate the new charge integration electronics, as well as to implement the gigabit transceiver (GBT) protocol used for the 4.8 Gbps (2.4 Gbps during operation) optical links. In the HCAL Barrel, the ngCCM is a three-card stack, with two control cards, each providing a communication link, and one clock card, for clock distribution. Figure 4.7 shows a stylized block diagram of an HB ngCCM.

Each control card has a multimode (MM) radiation hard and magnetic field tolerant optical transceiver, called a VTRx [58], and a Microsemi Igloo FPGA. The clock signal is decoded from the incoming data stream within the FPGA on the control card and is passed to the clock card through the board-to-board (B2B) connectors to be amplified and distributed via clock chips.

Each module undergoes two phases of testing. First, the module is assembled and experiences a standardized quality control to ensure that the module meets the design specifications. Second, the ngCCM joins the rest of the new front end in “burn-in”,




Figure 4.7: A block diagram of an HB ngCCM. Only one control card interfaces with the backend concurrently.

where the electronics are run in a detector-like teststand for two weeks to reveal any initial-operation failures.

Prior to assembly of the full module, each single card undergoes quality tests and preparation. First, general resistance measurements are taken to ensure there are no shorts. The control card PCBs then have the VTRx installed, and the FPGA is programmed with mixed HBHE firmware, since the control cards are interchangeable between the systems. The firmware differs between the primary and secondary cards, so equal numbers of both are programmed. The main difference between the firmware flavors is the behavior at power-on: a primary card defaults as the master control card and decodes the clock for distribution. The secondary card defaults to slave-mode operation, and serves passively. Single control card checks finish by setting hardware switches on the top of the card. This signals that the card will operate in HB mode and that the debug mode is enabled, to prepare for the full module tests. The clock cards face similar checks and then have a large t-shaped heat sink installed to support and cool the 2.5V and 3.3V DC/DC converters that are installed next. The 2.5V powers the VTRx, is stepped-down to 1.2V to power the FPGA, and serves as a reference voltage for the on-board analog-to-digital converters (ADCs). The 3.3V powers the clock chips, the ADCs, and several other active components. Once the DC/DC operation is verified, the full module is assembled, alternating the





Figure 4.8: All components of a ngCCM, partially disassembled. Module is constructed left to right. The clock card is already installed between the bottom and a middle layer of housing in the leftmost part of the image. Additional copper heat sinks, called cooling fins, are also installed over the VTRx. The blue thermal gap pads used are visible on the heat sinks.


boards with the support mechanics, as shown in Fig. 4.8. In addition to support and stabilization, the mechanics also provide thermal conductive cooling. Each PCB is in contact with some part of the aluminum housing, which in turn is in contact with the water-cooled aluminum RBx shell.

In the early days of testing, it was discovered that the VTRx was not operating at a temperature expected from the ngCCM design. While the roughly 45°C operating temperature is within the tolerance of the VTRx (which is qualified for safe operation up to 60°C), extended operation at elevated temperatures shortens the lifetime of electronics. To reduce the operating temperature, a copper heat sink, called a cooling fin, was added to the HB ngCCM mechanics. The cooling fin covers the optical electronics and the laser driver on the PCB (printed circuit board) of the VTRx, and connects these heat sources with the aluminum housing. Figure 4.9 shows an installed cooling fin and its temperature effects. With the addition of the heat sink, the VTRx operates nearly 15°cooler. This temperature study produced an important design difference between the HCAL Barrel and HCAL Endcap ngCCMs, which became important in the investigations featured in Section 4.3.

After assembling a full module, the ngCCM is checked to ensure production specifications are met. The tests and the teststand are designed to probe all aspects of the ngCCM's role in communication and control. The ngCCM monitors and distributes clock to the entire front end and provides bidirectional communication with the back end. A four-part ngCCM test stand was designed to probe these behaviors. The elements of the teststand were initially developed for HF ngCCM testing, but were re-purposed for HE and HB ngCCM quality control. A snapshot of the ngCCM teststand is shown in Fig. 4.10.



(a)



(b)

Figure 4.9: (a) Cooling fin installation on the VTRx from a view of the face of the control card. The module is partially disassembled to reveal fin. The cooling fin is in physical contact with the optical sub-assemblies of the VTRx via a 7 mm x 3 mm 0.1 in thick Bergquist thermal gap pad. The VTRx transmitter ASIC is also in contact with the larger plane of the cooling fin with a 12 mm x 12 mm 0.06 in thick Bergquist thermal gap pad (not visible). (b) Temperature measurements for the VTRx in an HB ngCCM. Measurements were taken with Capteur RTD PT100 PR PRO probes, and aluminum housing measurements were verified against internal ngCCM temperature sensors.




Figure 4.10: Photo of the ngCCM teststand in use with a fully assembled HB ngCCM with the face plate installed. The purple PCB is the front end emulator. The green PCB is the JTAG test board, currently in use. The debug interface is not connected. The blue optical fiber protruding from the ngCCM is connected to the primary control card, and the edge of the Glibv2 can be seen to the left side, on top of the power supply. Additional HB ngCCM can be seen.

The primary part of the teststand is a front end emulator. The emulator provides a mock-backplane to power the ngCCM and I2C buses to test the ability of the ngCCM to communicate via I2C across the backplane. The emulator also provides probe points for the distributed clock signals and LED indicators for the presence of counter signals. The emulator's LED thresholds are tuned to be sensitive to the ngCCM's output specifications. The second part of the teststand is a GLIBv2, a board with a Xilinx FPGA and an industry-standard optical transducer. The GLIBv2 provides a compatible backend system to optically communicate with the ngCCM. This exercises the optical link and the GBT veracity. The two final parts of the test stand are ancillary: a debug interface bypasses the optical control link for secondary checks and an additional board tests the ngCCM's ability to program FPGAs via the backplane. Without this ability, the firmware on the QIE cards would not be able to be updated.

Most of the tests are pre-programmed routines that require tester-input (power cycles to break communication, verification of resulting LED signals, etc). The most

common cause of module failure was poor quality clock distribution to the backplane, which was easily fixed by swapping the clock card. Replacing the problematic clock drivers would recover the clock card. Another show stopper (though never a final curtain call) was the Gigabit Transceiver (GBT) test, which probed the ability of the ngCCM to recover communication after interruption. The tests checked for the accumulation of pseudo-random bitstream (prbs) and Reed-Solomon decode (rsdecode) errors throughout power cycles and at various points in the full test procedure, to ensure the quality of the links. If these tests and others were successful, the ngCCM was passed to the “burn-in” stage of Phase 1 system testing.

At the peak of burn-in, thirteen RBx shells were populated with the new front end electronics and operated continuously for two weeks to flush out any beginning-of-life errors and failures. The burn-in teststand has the same backend controls as the official CMS system, allowing the new electronic suites to be tested under realistic conditions. In the end, around 100 HB ngCCMs were built and qualified as either installation ready, or installation-quality spares. The burn-in installation work informs the underground installation procedures, and is a prerequisite for joining the installation team.

4.2.3 HCAL Barrel Installation

As the electronics become available after their burn-in cycle, the next step is to install the new front end in the CMS detector. The front end electronics ring the outermost sections of the interior magnet volume, so the Surkov frame must be installed to allow for 360° access. The Surkov frame, shown in Fig. 4.11, is a specialized scaffolding that features a central region where a person can stand and smaller platforms along the edges of the magnet volume to allow access at varying heights. The frame is mounted from the support structure of CMS and allows for the entire endcap nose volume to be used as workspace. To access the Surkov frame, either a scaffolding tower is built from the floor of the cavern, or a passerelle for one person is built from level two within the cavern. The passerelle is ideal for when the yoke position requires the endcap to be flush with the edge of the magnet volume.

Due to the 360°-geometry of the HCAL Barrel, the ease of access to each RBx varies wildly, and in some cases presents both technical and physical challenges. Figure 4.12 shows a few different access positions. Figure 4.12a shows the access requirements of sectors 14 and 15, which are at 6 o’clock in the barrel. These sectors are the simplest from the technical point-of-view: the RBx normal is anti-aligned with gravity, and therefore the RBx is in its “upright” position. The challenge lies in the installer’s position: access is only achieved laying prone on the floor of the vacuum tank, with one’s arms outstretched overhead. Contrary to intuition, the simplest RBxs to physically access are sectors 1 and 10, which are geometrically at 3 and 9 o’clock. As





Figure 4.11: The left image shows the minus-side Surkov frame, connected via the passerelle. The right image was taken from the central platform inside the frame looking out of the frame, with the endcap nose in view.


shown in Fig. 4.12b, here the installer can be upright and work on the RBx; however, the RBx normal is perpendicular to gravity and the stabilization mechanics on the electronics receive the most stress as compared to other sectors.

Physical challenges strike each location differently, but the procedure was always the same. To begin, the Phase 0 electronics are removed. Next, the new Phase 1 backplanes and voltage modules are installed. Once supply voltage is confirmed, ngCCMs are installed, their optical fibers connected, and communication verified. In most sectors, these components are all hidden behind the tracker cable trays, and must be done by feel. Once the ngCCMs are in place, the calibration units and readout modules are installed, and a first round of commissioning is performed with the calibration unit fibers connected. Once communication has been confirmed with all active electronics, the read out modules are cleared for megatile connection.

The cable trays are routed through the middle of the RBx where there are no optical fibers. To accommodate for the new optical decoder unit mapping, some megatile fibers are swapped or replaced to meet minimum length requirements. Before megatile connection, we systematically order the megatile fibers to avoid tangles, and ease the connection process. Prior to each connection the megatile fiber face is cleaned and checked for fiber-breaks. Initially, connecting the tiles to two of the RMs took about



(a) Bottom sector access.



(b) Side sector access.



(c) 8 o'clock sector access.



(d) Top sector access.

Figure 4.12: In (a) Alexey Kalinin accesses a bottom sector. The image in (b) shows the author braced to access a side-sector in the barrel. The image was taken from the floor of the vacuum tank. The sector accessed in (c) can only be accessed from a crouch from a beam, and the top sectors (d) are hard due to fighting gravity, and because one cannot stand-up straight.

a half-day. With practice, an entire RBx can be connected in two hours, depending on the position.

The vestiges of the Phase 0 system provide the main challenges when connecting the scintillating tiles to the new read out electronics. The original fibers are retained as much as possible, since the active material is unchanged. The Phase 0 megatile fiber length is not optimal for the new readout module geometry. The fibers for the more-central layers now have excess length, and bend in a large arch when attached to the RM. If extreme care is not taken at connection, and in some cases, even in spite of this, closing the sectors with the original aluminum covers is challenging, since the fibers arch into contact with the cover. Figure 4.13a shows two properly connected RMs.

The original calibration system remains as well. The HCAL barrel and endcaps are both calibrated by running radioactive sources along tubes, called “source tubes”, inside the tile trays. Figure 4.13a also shows the sourcing tubes protruding from the megatile fiber bulk. This illustrates the primary challenge of megatile connection in the Phase 1 system. The sourcing tubes must be accessible. In a perfect world, the sourcing tubes would be routed in front of the calibration unit (right side of Fig. 4.13a, in the gap behind lime-green cables), and not in front of the RMs. This was the case in the Phase 0 system, but the calibration unit was moved in the Phase 1 system. Despite this, the sourcing campaign following megatile connection was successful, and few errors or damaged fibers were found. With the last RBx connected, we declared the upgrade complete, and transitioned into the operation and commissioning of the new system.

4.3 VTRx Investigations


The staged schedule of the HCAL upgrades offered several advantages, one of them being that the HCAL Endcap upgrade, which featured much of the same technology as used in the Barrel upgrade, was exercised fully during the 2018 data-taking period. Unfortunately, the CMS HCAL was the first to observe the optical-link loss associated with “RSSI drift” and attribute the loss to the VTRx, leading to a 1.5 year long investigation that ultimately impacted all of the LHC experiments and beyond [59].

4.3.1 Description and extent of the phenomenon

In August 2018, communication with a control link in one HE ngCCM was lost, resulting in the loss of data from this upgraded HCAL Endcap sector. Data taking resumed with a switch to the redundant control link, but the issue became a pattern when communication was lost for a second time in the opposite endcap during a



(a) Megatile fibers connected to RMs. (b) close-up of fiber-RM connection.



(c) Fiber organization from megatile-end connection.

Figure 4.13: Image (a) shows two connected RMs through the access hole. The thick black fibers are the megatile fibers. The black tubes with metal connectors are used for calibration. The light blue fibers are control fibers. The bright blue tubes and the red cable are the dry gas and high voltage, respectively. Image (b) shows a close-up of a partially connected RM. The Layer 0 connector (right, without fiber) has an opaque reflection, showing a damaged RM optic. Image (c) shows the organized megatile fibers from the megatile-connection end.




Figure 4.14: RSSI plot for all active links in the affected endcap over the lost module’s corresponding drift period. The HEM09 and HEP06 deviations are clear in comparison to the relative stability in the other modules.

commissioning run in 2019. Prior to the two instances of communication loss, the only indication of an unhealthy system was a phenomenon that became known as “RSSI Drift”, an observed gradual reduction in the Received Signal Strength Indicator (RSSI) over time, as shown in Fig. 4.14. The RSSI is an output of a current mirror of the PIN diode bias circuit within the optical receiver on the VTRx.

The RSSI current is proportional to the amount of light received at the front end. As received light decreases, the RSSI decreases as well. The average RSSI value in the endcap is $350\ \mu\text{A}$, with all modules in the $250\ \mu\text{A}$ to $450\ \mu\text{A}$ range, corresponding to approximately $-3\ \text{dBm}$ in received optical power. Prior to both communication loss events, the RSSI current decreased over time below the $10\ \mu\text{A}$ range, equivalent to a dark optical fiber. Immediately upon module extraction, the optical path between the backend and the ngCCM was checked, and no fiber damage was detected. With the most natural explanation excluded, the reduction of transmitted light to the front end, we began to investigate the interplay between the custom HCAL electronics and the VTRx.

Upon close inspection, forty-five percent of all HE links exhibited some level of RSSI instability, defined as an RSSI loss greater than $20\ \mu\text{A}$ from the initial value. Instabilities manifested in different patterns, but the two failures occurred after periods of exponential-like decrease in RSSI, with communication failing after the RSSI dropped below $10\ \mu\text{A}$. Most peculiar about the instabilities was the timescale of the drift. For the 2018 and 2019 communication failures, the time scales were four months and three days, respectively. The non-critical drift structures stabilize within four to twelve hours after powering.

Localizing the issue proved challenging due to the difficulty of reproducing the drift. Once modules with problematic links were removed from the CMS cavern to above-

ground teststands, the RSSI trends were nearly impossible to replicate. After the rest of the ngCCM had been ruled out, a teststand capable of operating and monitoring a single VTRx outside of the HCAL ngCCM hardware was built, and this setup replicated the drift features observed on-detector. With only one teststand available to capture the infrequent drift, we were unable to determine if drift was a property of a few faulty VTRxs or if the behavior was intrinsic to all VTRx modules.

By this time, the ALICE experiment had reported VTRx induced communication loss, corresponding with a drift in RSSI. A cross LHC working group was formed of representatives from CMS HCAL, CMS Electronics coordination, ALICE, and CERN-ESE who designed the VTRx. With CERN-ESE's stock of VTRxs and resident expertise, they replicated the drift en masse, and showed that given time, drift could arise in any VTRx.

4.3.2 VTRx Problem: HCAL Solution

The use of VTRxs in the CMS HCAL endcap is not unique within the HCAL sub-detectors. In addition to the 72 multi-mode (MM) VTRx links in HE, there are 144 identical control cards and VTRxs installed in the HCAL barrel. The ALICE collaboration discovered separately that single-mode (SM) VTRxs drift [60], implying that HF could also be affected. At the time of these studies, there were 16 single mode (SM) VTRxs in the HCAL forward (HF). The pervasiveness of the issue implied that the VTRxs in HB and HF should also experience drift. The HCAL barrel had been extensively operated on the primary link in commissioning runs, and during this period, no RSSI drift was observed. HF has been operating without RSSI drift since the VTRx installation in 2016.

The control electronics in HE and HB are both situated on-detector in water-cooled aluminum housings, with the water temperature held to around 18 °C. The HE and HB ngCCM share the same control card PCB design, making the primary difference between the ngCCMs the use of two clock cards in HE, as opposed to the single clock card in HB, and the addition of the cooling fin on the VTRx. The cooling fin was added only after HB production had begun, a year after the HE installation. The HF ngCCMs operate in an entirely different regime, installed on a mezzanine on a blade in an air-cooled rack off-detector, connected to the optical link via a pigtail within the body of the blade. The HF VTRxs operate around 38°C, measured from the VTRx receiver. It was not clear what HB and HF had in common to make them immune to the drift, but the cooling fin on the VTRxs in HB was a promising clue to the HE and HB difference.

Early in the investigations the cooling fin was suspected to mitigate the drift. As luck would have it, the first batch of HB ngCCMs to be sent to burn-in were built prior to the temperature studies and fin production, therefore they operated without cooling

fins. Monitoring data existed for this time. Three HB links were discovered to have drifted during this period. Out of those three HB ngCCMs, two had been reworked to install fins: one module was now installed on detector in HBP17, and had not drifted in its second burn-in or on detector, and the other was a spare. The third had been a prototype ngCCM to fill out burn-in tests, and its VTRx had been harvested and reused. We left HBP17 in place and quickly began testing HB ngCCM 60 (the spare that had drifted in burn-in without a fin, but had remained stable since the fin was installed). Drift could not be triggered in HB ngCCM 60. Without replication, we could not definitively say that the cooling fin had mitigated the drift.

Even though the HB ngCCMs that had been installed were operating stably, they had not been operated as long as HE, and the burn-in results without cooling fins remained a curiosity. The COVID-19 pandemic brought HB commissioning to a halt in early 2020. Because of this, the secondary links were not studied as judiciously as the primary. As we were slowly allowed to operate the detector again a more thorough survey of the secondary links in HB resumed. A single HB link with RSSI loss was finally observed. Only losing $10\ \mu\text{A}$ a day in RSSI, this “slow drifter” was only observed during the first two continuous weeks of detector operation after the COVID-19 lockdowns at CERN. Prior to that, we had only been able to power HB for one or two days at a time, which did not allow the RSSI to drift low enough to be considered a problem. Upon extraction from the detector, the suspect link was discovered to lack the cooling fin due to an assembly error. This fortuitous discovery, in light of our other indirect knowledge, indicated beyond a doubt that the cooling fins mitigated drift in HB. At this point, we proceeded to remove, rework to add a cooling fin, and retest all of the HE ngCCMs, following the same procedures outlined in Section 4.2.2.

4.3.3 Cooling Fin Properties

Even with our solution in hand, the root of the drift remained elusive. The most natural conclusion was that the VTRxs were operating too hot, and that the cooling fins reduced the temperature enough to stop the drift, but our HF data directly contradicted that conclusion. Since HE ngCCMs drifted at temperatures around $42\ ^\circ\text{C}$, the nominal HF operation around $38\ ^\circ\text{C}$ seemed too close to operate without drift. Furthermore, all previous temperature tests of the VTRx drift were either inconclusive or showed no effect. This discrepancy suggested that the cooling was not the end of the story, so we launched direct cooling fin investigations.

The cooling fins have three properties:

1. They provide conductive cooling
2. They are electrically grounded

3. They apply a mechanical pressure

We performed tests that isolated each of these properties. Around this time, HB ngCCM 60 had begun drifting reliably after several months without the cooling fin installed. Between each test, we checked to make sure the HB ngCCM was still drifting but monitoring operation without the fin for at least one-hour. After the check, we would let it rest with the optical fiber unplugged from the VTRx. The most straightforward and least invasive test was to check the grounding properties. To do this, we electrically isolated the cooling fin from the grounded ngCCM housing using Kapton tape. The cooling fin continued to mitigate drift. For the next iteration, we thermally isolated the fin by replacing the thermal gap pad with Armaflex. In this scenario, the drift proceeded as if the fin was not installed. Since the Armaflex is not as firm as the thermal gap pad, we wanted to remove any doubt that the mechanical difference was the source of the drift, so we replaced the gap pad once more, but this time with rubber. In this scenario, we more closely reproduced the mechanical properties of the properly installed fin. This produced a drift structure halfway between mitigation and full drift. It turns out that rubber is 10 times more heat conductive than Armaflex, implicating the cooling property of the fin.

Temperature was mitigating the drift, but the discrepancy between the operating temperatures of HE/HB and HF hinted that the absolute temperature is not the governing factor. Since these tests had been done with an entire RBx, the exact temperature landscape of the VTRx in these scenarios remained a mystery. A smaller, single VTRx teststand with better temperature control was constructed, and a VTRx with reliable drift lent to us by ALICE was installed. This teststand included temperature measurement and control for the incoming fiber ferrule, the receiver canister, and other VTRx and optical plant components.

4.3.4 Problem origin

The modularity of the single-VTRx teststand allowed for the observation of the condensation of a material on the face of the incoming fiber to the VTRx. This was the origin of the drift. Over time, material outgassing from within the VTRx would condense on the fiber face, obscuring the transmission of light. Immediately after this discovery, the multi-VTRx teststands from other collaborators confirmed the presence of the outgassed material in their drifting VTRxs. The results of the outgassing are pictured in Fig. 4.15. What originally appeared to be a reduction of light to the front end was indeed that, but the source of the obstruction was the front end itself! The resulting communication losses would occur when so much material condensed on the fiber face that the fiber appears to go dark. The material was extremely volatile, and would evaporate over time, complicating the efforts to identify the source of the problem.

Two processes determine whether the VTRx will drift: the outgassing process itself, and the condensation reaction. This offers two solutions to the RSSI-correlated communication loss. First, one could inhibit the outgassing. The CERN-ESE team has developed a 500 hour bake-out procedure for extant VTRxs, but due to the nature of the problem, only a redesign of the VTRx mechanics would be a complete solution. The second solution, and the one employed by CMS, is to stop the condensation. Figure 4.16 shows the impact of the temperature differential between a VTRx and the incoming fiber ferrule on the drift. In this case, by maintaining the temperature difference between the incoming fibers and the VTRx receiver to around 10°C , the RSSI stabilizes. Minimizing the temperature difference is preferred.




Figure 4.15: Image of fiber face after removal from drifting VTRx (left). Outgassed material appears as a white substance causing the perceived texture on the fiber face, and obscuring the central part of the fiber where light is transmitted(see red circle). Image of clean fiber for comparison (right).Red circled region is clear.

The nature of condensation makes it the only mechanism consistent with our various observations throughout 2020. We speculate that the lack of reproducibility of the drift in the initial studies was likely due to the season of the year. The measurements were done in the summer, and the ambient temperature of the fibers in the laboratory were higher than the fibers in the CMS cavern, and probably high enough to stop condensation. This explained one particularly notable observation. One morning we had powered-on all of the teststands after 24 hours of rest, and every link had drift, for the first time in the laboratory. That was never reproduced. In hindsight, meteorological data showed that the night prior had been one of the cooler ones in August 2020.




Figure 4.16: Mapping of operational phase space taken from measurements of a single, illustrative VTRx experiencing RSSI drift. This shows how the temperature difference between the fiber ferrule and the VTRx receiver influences the condensation of the outgassed material. The z -axis is the derivative of the RSSI, with a negative trend indicating the RSSI was decreasing, and a positive trend indicating that the RSSI was increasing. The trend in the RSSI reveals if the temperature settings either encourage condensation, or are unfavorable to condensation. Measurements are not available for the gray region.

4.3.5 Final Remarks

It was clear that the communication losses suffered in the HCAL upgrade were due to outgassing in the VTRx. The outgassed material would condense on the face of the fiber, blocking transmitted light. We mitigated this effect without direct intervention to the VTRx by minimizing the difference in temperature between the incoming fiber and the VTRx receiver with the addition of a heat sink on the VTRx. The CERN-ESE team has since concluded that the outgassing is due to a UV adhesive used during assembly and its reaction to the metal of the receiver canister.

Since the addition of a heat sink is a mitigation tactic, the CMS HCAL maintains the risk of data loss if temperature control is lost. Despite this, with our solution, we believe we are capable of stable operation for the remainder of the CMS HCAL Barrel's operation. After the rework, no drift has been observed in HE. Other VTRx users have started implementing heat sinks inspired by our design when other solutions are not viable.

Chapter 5

Particle Reconstruction

CMS is a *particle flow* detector and is designed to exploit the complementary information between the subdetectors to determine particle kinematics and identification. The subdetectors' measurements are combined to build particles in a process known as *reconstruction*. Muons transverse the entire detector and leave hits in the tracker, small deposits in the calorimeters, and exit with hits in the outer muon systems. Both electrons and photons deposit most of their energy in the ECAL, but electron trajectories in the tracker complement this measurement, and provide electron/photon differentiation. Jets, or the collimated multi-particle decays of hadrons, leave signatures across all of the subdetectors: charged hadrons have tracks in the tracker, ECAL deposits, and HCAL deposits; neutral hadrons deposit energy in the HCAL; and heavy-flavor quark decays can even have muon system hits. Particles that primarily interact weakly, like neutrinos and potential BSM particles, exit the detector without interacting, but their presence can be inferred from the measured momentum imbalance in the transverse plane, called $\vec{p}_T^{\text{miss}}$. Figure 5.1 shows a simplified cartoon of the different particle interactions in a transverse view of the CMS detector.

Section 5.1 introduces the Particle Flow (PF) algorithm used in CMS. Particle Flow serves as the basis for the dominant event reconstruction in CMS. Sections 5.2 and 5.3 discuss muon and electron reconstruction, respectively. Jets are built out of collections of reconstructed particles, and are covered in Section 5.4. Finally, once all of the visible particles are reconstructed, composite objects like $\vec{p}_T^{\text{miss}}$ can be built. Section 5.5 describes $\vec{p}_T^{\text{miss}}$ derivation.

CMS is sensitive to all stable and quasi-stable particles that are produced in LHC collisions. Notably absent in the above enumeration are photons and tau leptons. Though they are not used in this analysis, details of their reconstruction can be found in Refs [61] and [62], respectively.

5.1 The Particle Flow Algorithm

The *particle flow* (PF) algorithm forms the building-blocks of the particle reconstruction in CMS [63]. For charged particles, the PF algorithm forms particle trajectories by matching hits in the silicon tracker. These trajectories are known as *tracks*.




Figure 5.1: Transverse section of CMS illustrating locations of different particle interactions. [63]

From the calorimeter deposits, the algorithm forms *clusters* of energy. These clusters are particularly important for the identification of neutral particles, and when used in combination with tracking, augment the reconstruction of electrons and charged hadrons in jets. Finally, the algorithm *links* tracks and clusters to form *blocks* that serve as the basic particle identification (ID).

Tracks are found in the PF algorithm through an *iterative tracking* process. Iterative tracking is the successive application of the combinatorial track finder (CTF) [64]. The CTF consists of three main steps: seed generation, trajectory-building, and fitting to extract the particle kinematics. The “seed” of the track is the initial detector signature that provides the first estimates of the track’s characteristics. In the first few iterations, three close-in-space pixel hits serve as the track seed; in the final iteration, muon chamber hits seed the tracks. The track building and fitting are done simultaneously with a method based on the Kalman Filtering (KF) formalism [65]. The KF formalism, when adapted to particle physics track building, builds a track by projecting the track already constructed (the seed, in the case of the initial steps) and some expected uncertainty onto the next tracking-detector layer. Hits on that

layer are then combined into the track, and if the overall uncertainty on the entire track is minimized, the hit is kept in the track, and the process continues to the next layer. Table 5.1 summarizes the tracks targeted in each of the ten iterations. After each step, the hits associated with the selected tracks are masked to reduce the combinatorial complexity of the next iteration. Beyond the seed configuration, each iteration has specific track fit χ^2 and primary-vertex compatibility requirements that further depend on the track's p_T , $|\eta|$, and number of tracker hits. Once tracks are made, they can be connected together to form the particles used in CMS analyses.

The second component of the PF algorithm is the generation of calorimeter clusters. Clusters are formed separately for each of the barrel and endcap portions of the ECAL and HCAL. Like tracks, clusters initiate around seeds, where seeds are defined as calorimeter cells with an energy deposition above some tunable threshold, and greater than the surrounding cells. The initial cluster is formed by gathering the neighboring cells, then *topological clusters* are created by adding cells that share at least a corner with a cell in the initial cluster and have an energy deposit greater than twice the detector noise. Table 5.2 summarizes the clustering parameters. Each topological cluster is assumed to arise from N Gaussian energy deposits (seeds) with some amplitude A_i , width σ (specific and predetermined for each calorimeter), and position in (η_i, ϕ_i) . An expectation-maximization algorithm performed in two steps finds the energy-fraction and positions of each of these clusters, with the seed values providing the starting parameters. The first step determines the energy fraction of each Gaussian deposit, and the second step determines the deposits' locations from an analytical maximum-likelihood fit. The resulting clusters serve as the foundational clusters for calorimeter objects in CMS.

Finally, the PF algorithm connects tracks and clusters across different subdetectors via the *link algorithm*. Elements are considered for linking if they are nearest-neighbors in (η, ϕ) . To link tracks to calorimeter clusters, the track is extrapolated in η and ϕ to the ECAL and HCAL. If this extrapolation intersects a cluster, the cluster and the track are linked. Calorimeter clusters are linked if the entire cluster in the innermost calorimeter is fully enveloped in the outer calorimeter's cluster. The link algorithm is limited by the detector granularity and the total particle flux. As one moves deeper into the detectors, nuclear interactions that change the trajectories of particles become more likely, and reduce the probability of linking all of the deposits originating from a particular particle. The final output of the linking algorithm are *PF blocks*

Each block typically contains information about several separate particles. Muons, which are described in Section 5.2, are reconstructed first, and their corresponding components are removed from the blocks. Electrons and photons are built next. Afterwards, charged and neutral hadrons are built from the remaining information.

Iteration Order	Step Name	Targeted Track Type
1	InitialStep	prompt, high p_T
2	DetachedTriplet	b hadron decays, $R \lesssim 5$ cm
3	LowPtTriplet	prompt, low p_T
4	PixelPair	recover high p_T
5	MixedTriplet	displaced, $R \lesssim 7$ cm
6	PixelLess	very displaced, $R \lesssim 25$ cm
7	TobTec	very displaced, $R \lesssim 60$ cm
8	JetCoreRegional	inside high p_T jets
9	MuonSeededInOut	muons
10	MuonSeededOutIn	muons

Table 5.1: Iterative tracking targets. Naming loosely corresponds to the seeding method used in each iteration. [63].

	ECAL Barrel	ECAL Endcap	HCAL Barrel	HCAL Endcap
Seed E Threshold (MeV)	230	600	800	1100
# of neighbor cells	8	8	4	4
Cell E Threshold (MeV)	80	300	800	800
Seed E_T Threshold (MeV)	0	150	0	0
Gaussian width σ (cm)	1.5	1.5	10.0	10.0

Table 5.2: Summary of the clustering parameters. The neighboring cells are defined either as those cells sharing an edge with the seed cell (for a total of four neighbors), or additionally sharing corners (for a total of eight neighbors). The ECAL endcap has an additional requirement on the transverse energy due to a non-uniformity in the noise dependent on θ . Quoted values are determined from optimizations based photon, jet, π^0 , and K_L^0 simulation.

5.2 Muons

Muon trajectories can be reconstructed both in the tracker and in the muon systems [66]. Hits in the CSCs, DTs, and RPCs are reconstructed to form *standalone muon tracks* using a Kalman-filter technique. *Tracker muons* have tracker tracks and at least one geometrically-matched muon system hit. *Global muons* extend the track-building from the standalone tracks to include tracker hits, in an “outside-in” Kalman-filtering technique. These three tracking scenarios correspond with the last two iterations of tracking in Table 5.1. When detector services block the muons systems, or when the muons have low- p_T , the tracker muons are the most performant. Global muons have better high- p_T muon reconstruction due to the greater deflection at large radii. Muon momentum and sign are extracted from the fits to the muon’s tracks. The sign of the muon’s charge corresponds with the sign on the trajectory’s curvature.

Different muon IDs are built on-top of these track types by tuning the following parameters of reconstruction:

1. track fit χ^2
2. Number of hits in the different subsystems
3. Degrees of matching between tracker and global muon tracks
4. Impact parameter

The values for these parameters are tailored to the target physics analysis. The details of the identification used in this analysis are given in Section 6.4.2.

5.3 Electrons

Electrons are primarily reconstructed in the ECAL and the tracker. Electrons emit bremsstrahlung photons as they move through the tracker material, and the emitted photons interact with the ECAL tangentially along the electron’s curved path in ϕ . Electrons are reconstructed as *superclusters* in ECAL, with calorimeter clusters grouped together at a fixed η , but in an extended region in ϕ to capture the photons emitted along the electron’s path. The ratio of energy reconstructed in ECAL and the momentum extracted from the track-fit is required to be close to one. Such a reconstruction method is termed *ECAL-based*, where the seeding comes from the ECAL deposits.

This method fails when electrons are non-isolated, or when the electrons are low in energy. In these cases a *tracker-based* reconstruction is used. For soft electrons, the

entire track can be reconstructed and directly linked to ECAL deposits; however, if high-energy photons are emitted, the track can dramatically change, and the uncertainty on the fit can be large. Tracks with fewer tracker hits and low χ^2 values are refit using a Gaussian-sum filter (GSF) [67]. The GSF is a generalization of the KF. The KF formalism assumes that the distributions of the state vectors used in the fit are single Gaussians, while GSF assumes that they are weighted sums of Gaussians, giving more flexibility. The χ^2 of the GSF track, the ratio of the GSF track χ^2 to the KF track fit χ^2 , the number of hits, the energy lost along the track, and the distance between the track extrapolation and the closet ECAL cluster are all then fed into a boosted-decision tree to identify the track as an electron.

The two tracking-collections, the ECAL-based and tracker-based, are then combined together into a single electron candidate collection, and are all refit with a more robust and computationally intense GSF. This final collection provides the electron tracking. Tracks are linked to calorimeter clusters as described in Section 5.1.

Electrons that originate at the primary vertex are identified using additional criteria [61]. These “prompt” electrons have isolation and shower-structure variable requirements applied to differentiate them from electrons within jets. The two primary shower-structure variables are the ratio of hadronic to electromagnetic energy (H/E) and the $\sigma_{i\eta i\eta}$, described below. The H/E is the ratio between the energy deposited in the HCAL found in a $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.15$ cone around the electron candidate direction and the reconstructed electron energy in the ECAL. Three real processes coordinate HCAL deposits with real, prompt electrons:

1. HCAL Noise
2. Pileup
3. Electrons/photons passing through gaps in ECAL

To model these different contributions, the H/E selection criteria takes the form

$$H < X + Y\rho + JE, \tag{5.1}$$

where X and Y are noise and pile-up terms, respectively, ρ is the median transverse energy density per area of the event, and J is a scaling term for high-energy electrons and photons. The J captures how the dominant contributions change with electron or photon energy.

Electrons and photons from the primary vertex tend to have more spatially concentrated energy deposits in the ECAL than electrons and photons originating from

hadrons. The $\sigma_{i\eta i\eta}$ is the second moment of the log-weighted distribution of crystal energies in η of the 5×5 crystals surrounding the most energetic crystal in the supercluster,

$$\sigma_{i\eta i\eta} = \sqrt{\frac{\sum_i^{5 \times 5} w_i (\eta_i - \bar{\eta}_{5 \times 5})^2}{\sum_i^{5 \times 5} w_i}}, \quad (5.2)$$

where η_i is the pseudorapidity of the i th crystal, $\bar{\eta}_{5 \times 5}$ is the mean position in pseudorapidity, and the weight w_i is given as $w_i = \max(0, 4.7 + \ln(E_i/E_{5 \times 5}))$. This requirement also reduces the impact of ECAL noise.

Another powerful parameter for electron identification is isolation. The isolation variables are sums of the transverse momentum of the charged hadrons, neutral hadrons, and photons in a cone of $\Delta R = 0.3$ around the electron direction. The isolation cut applied is dependent on the transverse energy of the electron in question, since the energy spread of the electron candidate is also dependent on its energy. Pile-up corrections are applied and are energy-dependent.

Tracker quantities are also used to identify prompt electrons. Angular variables like $|\phi_{\text{supercluster}} - \phi_{\text{track}}|$ and $|\eta_{\text{seed}} - \eta_{\text{track}}|$ check agreement between calorimeter clusters and tracks. Electron momenta are measured both in the ECAL and in the tracker, so variables like $|1/E - 1/p|$ check the two different measurements' agreement, where p is the momentum extracted from the tracker measurement, and E is the ECAL measurement. Electrons originating from the primary vertex should have hits in the innermost tracking layers, while electrons from photon-conversion (with the photon coming from the primary vertex) will leave hits only in the deeper layers. Requirements on inner-hits help distinguish electrons from photons.

Electron energy is estimated by combining the ECAL and tracker measurements,

$$E_{\text{combined}}^{\text{reco}} = \frac{E_{\text{ECAL}}/\sigma_E^2 + p_{\text{tracker}}/\sigma_p^2}{1/\sigma_E^2 + 1/\sigma_p^2}, \quad (5.3)$$

where σ_E and σ_p are the energy resolution of the supercluster and the momentum resolution of the electron tracking algorithm, respectively. The energy regressions and corrections are given in more detail in [61]. The sign of the electric charge carried by the electron or positron is revealed by the sign of the curvature on the GSF track.

5.4 Jets

Due to color confinement, quarks and gluons hadronize. The quark bound states then interact and decay, producing collimated sprays of particles called jets. To reconstruct the initial particle, the jet reconstruction must capture all of the daughter particles. The PF algorithm reconstructs all of the final state particles in the event, and these particles are clustered into jets using the *anti- k_T* algorithm [63, 68]. The anti- k_T algorithm clusters particles according to the following distance measure:

$$\begin{aligned} d_{ij} &= \min(p_{Ti}^{-2}, p_{Tj}^{-2}) \frac{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}{R^2}, \\ d_i &= p_{Ti}^{-2}, \end{aligned} \tag{5.4}$$

where d_{ij} represents a distance variable between particles i and j , y is the rapidity, p_{Ti} is the transverse momentum, and R is the maximal radius parameter in the pseudorapidity-azimuth plane. To decide whether to cluster two particles, one first defines the smaller of the d_{ij} and d_i to be d_{min} [69]. If the smaller of the two is d_{ij} , the particles are clustered together, their transverse momenta are summed, and the process begins again. If the d_i is the smaller of the two, the particles are not clustered together. The anti- k_T algorithm clusters around the hardest particles first, and continues until there are no particles left. In CMS, we reconstruct jets using both $R = 0.4$ and $R = 0.8$. These two types of jets are called “AK” (anti- k_T) “4” ($R = 0.4$) and “AK8” jets, respectively.

We use pileup per particle identification (PUPPI) to mitigate the effects of particles that do not originate from the primary vertex in the clustered jet [70, 71]. The PUPPI algorithm assigns each jet constituent four-vector a weight corresponding to how “pileup-like” the particle is based on its relationship in space to the surrounding particles. Particles from hard scatters are more likely to be collinear, while pile-up will be diffuse in space because these particles originate from multiple vertices. The charged particle distributions from both the primary and pileup vertices help to build the shape variables used to derive the weight. This assumes that neutral particles from the primary vertex will cluster near charged particles from the primary vertex and vice versa for the neutral contributions due to pile-up vertices.

When a heavy particle, like an intermediate vector boson or a top quark, has a large momentum, the resulting decay products are collimated. For hadronic decay channels, this results in overlapping jets. Instead of attempting to resolve daughter particles as separate, small radius AK4 jets, we can cluster all of the daughter particles into a single large-radius AK8 jet. This large radius jet is often called a fat jet. This object is directly interpreted as the parent particle, eliminating combinatoric complications in the reconstruction, and allowing more sophisticated jet-algorithms to extract relevant

physical parameters. The most important fat jet parameter used in this analysis is the soft drop mass [72]. The soft drop algorithm pairwise declusters a jet, and removes soft constituents unless

$$\frac{\min(p_{T1}, p_{T2})}{p_{T1} + p_{T2}} > z_{\text{cut}} \left(\frac{\Delta R_{12}}{R_0} \right)^\beta, \quad (5.5)$$

where z_{cut} is the soft drop threshold, R_0 is the jet radius, ΔR_{12} is the separation, and β is the angular exponent. Soft drop (SD) is a generalization of generic jet-grooming algorithms, with the goal to reduce wide-angle and soft contributions that obfuscate the parent particles mass.

The composite nature of jets provides challenges in reconstructing the true energy of the parent particle. Calorimeter response variation and pileup obscure the true measurement of jet momentum and energy. To account for this, Jet Energy Corrections (JECs) are derived to correct the Jet Energy Scale (JES) and Jet Energy Resolution (JER) of jets in simulation and to calibrate jets in data [73, 74]. Data calibrations account for pileup and residual differences with respect to simulation. The JECs are applied based on the η and p_T of the jet and are derived specifically for the jet radius and pileup mitigation algorithm. The JES corrections are derived in four steps, with each step informing the next:

1. Pileup offset
2. Detector response corrections taken from simulation
3. Corrections for differences between data and simulation
4. Optional corrections for jet flavor

The pileup offset is derived by taking the difference in p_T from simulated samples with and without pileup. This difference is then subtracted from data and simulation. The detector response corrections account for the differences in reconstructed and true energy, derived by comparing the particle level output to the full detector simulation. Finally, small corrections to match simulation to data are derived to account for the changes in the detector over time and to reflect the true detector performance. The flavor corrections are applied to reweight jets to account for differences in calorimeter response to hadron flavor, but these corrections are not applied in this analysis.

Scale factors are derived to match the JER in data and simulation. Jets in simulation are smeared to have the same resolution as jets in data. Jet smearing is performed in two ways [75]. When a simulated jet can be matched to a particle-level jet, the *scaling method* is used. In the scaling method, the jet four-vector is corrected by a factor of

$$c_{\text{JER}} = 1 + (sf_{\text{JER}} - 1) \frac{p_{\text{T}} - p_{\text{T}}^{\text{ptcl}}}{p_{\text{T}}}, \quad (5.6)$$

where sf_{JER} is the derived scale factor, p_{T} is the transverse momentum of the reconstructed jet, and $p_{\text{T}}^{\text{ptcl}}$ is the transverse momentum of the particle-level jet. When the reconstructed jet cannot be matched to a particle level jet, the corrected factor is derived via *stochastic smearing* giving

$$c_{\text{JER}} = 1 + \mathcal{N}(0, \sigma_{\text{JER}}) \sqrt{\max(sf_{\text{JER}}^2 - 1, 0)}, \quad (5.7)$$

where σ_{JER} is the p_{T} resolution in simulation, sf_{JER} is the derived scale factor, and $\mathcal{N}(0, \sigma_{\text{JER}})$ is a random number sampled from a normal distribution centered at 0 and a standard deviation of σ_{JER} .

5.5 Missing Transverse Momentum

The presence of weakly interacting particles in the final state is inferred from the momentum imbalance in the transverse plane of the detector. This missing transverse momentum, $\vec{p}_{\text{T}}^{\text{miss}}$, is computed as the negative vector sum of all of the PF candidates originating from the event's primary vertex [76]. PF candidates clustered into jets receive JECs, and these corrections are propagated to the $\vec{p}_{\text{T}}^{\text{miss}}$. This corrected $\vec{p}_{\text{T}}^{\text{miss}}$, referred to as *type-1 corrected*, is given as

$$\vec{p}_{\text{T}}^{\text{miss}} = \vec{p}_{\text{T}}^{\text{miss,raw}} - \sum_{\text{jets}} (\vec{p}_{\text{T,jet}}^{\text{corr}} - \vec{p}_{\text{T,jet}}), \quad (5.8)$$

where the jets are AK4 jets. These AK4 jets are not subject to a lepton veto. To accommodate this, jets with an electromagnetic energy fraction greater than 0.9 are excluded from the sum, and the four-vectors of standalone muons are subtracted from the four momentum of the overlapping jet (if applicable). The JECs are applied based on this modified jet-momentum for the correction in eq. 5.8.

Detector defects and particle mismeasurement contribute to “fake” $\vec{p}_{\text{T}}^{\text{miss}}$. To reduce the effects of known detector issues with $\vec{p}_{\text{T}}^{\text{miss}}$ reconstruction in certain data-taking periods, we apply the filters in Table 5.3.

	2016	2017	2018
Primary Vertex Filter	yes	yes	yes
Beam Halo Filter	yes	yes	yes
HBHE Noise Filter	yes	yes	yes
HBHE Iso Noise Filter	yes	yes	yes
ECAL TP Filter	yes	yes	yes
Bad PF Muon Filter	yes	yes	yes
EE BadSC Noise Filter	data only	data only	data only
ECAL Bad Calibration Filter Update	n/a	yes	yes

Table 5.3: Applied p_T^{miss} filters. Filters are applied to both simulation and data unless otherwise specified. Filters handle known detector or reconstruction effects. The HCAL Phase 1 upgrade should remove the need for the HCAL filters.

Chapter 6

Strategy and Event Selections

The baryonic leptophobic Z' model introduced in Chapter 2 offers a rich phenomenology ripe for analytical innovation. Cascade decays to invisible particles, like the anomalous in the aforementioned model, are not new in collider physics – but the *resonant* production of a Z' initiating the cascade is a novel feature. We can search for a mass peak if we can reconstruct the resonance. A resonance search offers a straightforward model interpretation and additional constraints to suppress non-resonant background processes.

In the boosted regime, the Z to dimuon is easily captured by the CMS trigger and jet-substructure can be used to discriminate the $H \rightarrow b\bar{b}$ against the SM backgrounds. Final states with such distinctive signatures are more tractable than the current dijet phase-space explored in leptophobic Z' searches. This opens a potentially more-sensitive region to exploration.

This thesis presents the first CMS search for a leptophobic Z' decaying to anomalous, as well as the first use of Recursive Jigsaw Reconstruction in CMS: the framework used to reconstruct the Z' peak. Section 6.1 introduces the phenomenology of the Z' final state we seek. Sections 6.2 and 6.3 detail the general analysis strategy and the final state observable. Section 6.4 goes into the specific physics objects use beyond the descriptions in Chapter 5 and how we reconstruct the neutral final state SM bosons. Finally, Section 6.5 summarizes the basic kinematic selections in this analysis.

6.1 Model Phenomenology and Final State

Ignoring the dijet decay channel of the leptophobic Z' introduced in Section 2.2.2, one is left with three decay channels. Assuming $m_{E_S} > m_{E_D} > m_{N_D} > m_{N_S}$ as the mass hierarchy, the following decays are allowed:

$$Z' \rightarrow E_S^+ \bar{E}_S^- \rightarrow W^+ W^- N_S \bar{N}_S + n(Z/H), n = 2, 3, 4, \quad (6.1)$$

$$Z' \rightarrow E_D^+ \bar{E}_D^- \rightarrow W^+ W^- N_S \bar{N}_S + n(Z/H), n = 1, 2, \quad (6.2)$$

and

$$Z' \rightarrow N_D \bar{N}_D \rightarrow N_S \bar{N}_S + nZ + mH, n, m = 0, 1, \text{ or } 2. \quad (6.3)$$

Example Feynman diagrams of these decays are given in Fig. 2.7. This analysis focuses on the decay chain in expression 6.3 because it has the most tractable final state. The charged-current decays can be used in future extensions of this search. The final state with one Z going to $\mu\mu$ and an H going to $b\bar{b}$ is chosen. This choice prioritizes the higher mass N_D phase-space. The dimuon decay has good background rejection and capitalizes on the CMS Experiment’s precise muon p_T resolution and efficient muon triggers. The $H \rightarrow b\bar{b}$ captures the largest H branching fraction.

The mass-splitting in a Z' -to-anomalon cascade determines the final state topology. The left figure in Fig. 6.1 illustrates the decay of a heavy $Z' \approx O(\text{TeV})$ with a daughter m_{N_D} of roughly half of the mass of the Z' . With both of the N_D particles produced almost at rest, the Z and H are produced with little correlation, and in some cases, might overlap. The isotropic decay maximizes the impact of the p_T^{miss} from the two N_S particles. Larger mass differences between the N_D and the Z' produce the opposite topology, illustrated in the right figure in Fig. 6.1. The final state is back-to-back and reduces the measured p_T^{miss} .

When the Z' mass is of the order of a TeV, the final state particles have high momentum, or are considered “boosted”. These boosted boson decays collimate the daughter particles. For the case of the $H \rightarrow b\bar{b}$, both daughter b-jets can be reconstructed as one large-radius fat jet. This large-radius jet improves the mass-resolution of the H reconstruction and capitalizes on jet substructure to identify the Higgs boson candidates over the other SM backgrounds. Section 6.4.4 describes the use of jets in this analysis.

6.2 Analysis Strategy

This analysis uses a data sample recorded by the CMS detector during Run 2 of the LHC with a total integrated luminosity of 137.6 fb^{-1} . The total integrated luminosity per-year is summarized in Table 6.1. Each year is broken down into different data-taking eras, corresponding to different detector configurations, and in some cases, warranting specific treatment.

The CMS data sets are defined by the trigger menu. The boosted dimuon Z in the final state aligns with the triggers in the SingleMuon data set. Each passing event must include a trigger-object corresponding to a muon with a $p_T > 50 \text{ GeV}$, or pass one of the more-restrictive back-up triggers. These triggers are included to improve the overall event selection efficiency. The complete trigger scheme for this analysis




Figure 6.1: Cartoon to illustrate two general classes of the model’s decay topologies. The left figure illustrates the scenario of nearly-at-rest production of heavy N_D particles. The right figure illustrates boosted N_D production, resulting in back-to-back signatures. This analysis prioritizes the former configuration where the SM backgrounds are strongly suppressed.

is given in Table 6.2. This analysis also uses the SingleElectron data set to build a background control region, which will be discussed in detail in Chapter 7.

Monte Carlo simulation (MC) assists the background estimation in this analysis, and models the Z' signal. Details of the MC samples used are included in Table 6.3. Signal samples are generated at leading order with **MADGRAPH5_v2.6.5** with a $g_{Z'}$ coupling value of 0.4. The cross section scales with the square of the coupling, so extracted limits can be interpreted in terms of different $g_{Z'}$ values. The cascade decay has three mass parameters: $m_{Z'}$, m_{N_D} , and m_{N_S} . Since the final state kinematics are dominated by the Z' and N_D masses, we explore two scenarios of m_{N_S} : a heavy $m_{N_S} = 200$ GeV and a light $m_{N_S} = 1$ GeV. Within these two scenarios, a mass scan in m_{N_D} and $m_{Z'}$ is performed. The theoretical cross sections for the $m_{N_S} = 200$ GeV scenario are shown in fig. 6.2. Figure 6.3 shows the explored signal grids.

Year	Luminosity	Uncertainty in percent
2016	36.31 fb ⁻¹	1.2 %
2017	41.53 fb ⁻¹	2.3 %
2018	59.74 fb ⁻¹	2.5 %
2016+2017+2018	137.6 fb ⁻¹	1.6 %

Table 6.1: Luminosities and uncertainties in Run 2 Collision Data recorded by the CMS Experiment.

2016	Muon	HLT_Mu50_v* or HLT_TkMu50_v*
2016	Electron	HLT_Ele27_WPTight_Gsf_v* or HLT_Ele115_CaloIdVT_GsfTrkIdT_v* or HLT_Photon175_v*
2017	Muon	HLT_Mu50_v* or HLT_TkMu100_v*
2017	Electron	HLT_Ele35_WPTight_Gsf_v* or HLT_Ele115_CaloIdVT_GsfTrkIdT_v* or HLT_Photon200_v*
2018	Muon	HLT_Mu55_v* or HLT_TkMu100_v*
2018	Electron	HLT_Ele32_WPTight_Gsf_v* or HLT_Ele115_CaloIdVT_GsfTrkIdT_v* or HLT_Photon200_v*

Table 6.2: Trigger paths used in Run 2 Collision Data.




Figure 6.2: Theoretical cross sections for the $m_{N_S} = 200$ GeV signal scenario with $g_{Z'} = 0.4$ and assuming a 100% branching fraction. The cross section peaks between $m_{Z'}$ values of 3500 GeV and 4000 GeV.




Figure 6.3: Grid of mass points generated for signal production.

Background MC Samples	$\sigma \times \mathcal{B}$ (pb)
2016 Samples:	
ZZTo2Q2Nu_TuneCP5_13TeV_amcatnloFXFX_madspin_pythia8	6.331
WZTo2Q2Nu_TuneCP5_13TeV_amcatnloFXFX_madspin_pythia8	3.688
DYJetsToLL_M-50_HT-100to200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	147.4
DYJetsToLL_M-50_HT-100to200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8_ext1	
DYJetsToLL_M-50_HT-200to400_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	40.99
DYJetsToLL_M-50_HT-200to400_TuneCUETP8M1_13TeV-madgraphMLM-pythia8_ext1	
DYJetsToLL_M-50_HT-400to600_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	5.678
DYJetsToLL_M-50_HT-400to600_TuneCUETP8M1_13TeV-madgraphMLM-pythia8_ext1	
DYJetsToLL_M-50_HT-600to800_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	1.368
DYJetsToLL_M-50_HT-800to1200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	0.6304
DYJetsToLL_M-50_HT-1200to2500_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	0.1514
DYJetsToLL_M-50_HT-2500toInf_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	0.003565
TTTo2L2Nu_TuneCP5_PSweights_13TeV-powheg-pythia8	72.1
TTToSemiLeptonic_TuneCP5_PSweights_13TeV-powheg-pythia8	314.0
TTToHadronic_TuneCP5_PSweights_13TeV-powheg-pythia8	301.0
2017 Samples:	
ZZTo2Q2Nu_TuneCP5_13TeV_amcatnloFXFX_madspin_pythia8	6.331
WZTo2Q2Nu_TuneCP5_13TeV_amcatnloFXFX_madspin_pythia8	3.688
DYJetsToLL_M-50_HT-100to200_TuneCP5_13TeV-madgraphMLM-pythia8	147.4
DYJetsToLL_M-50_HT-100to200_TuneCP5_13TeV-madgraphMLM-pythia8_ext1	
DYJetsToLL_M-50_HT-200to400_TuneCP5_13TeV-madgraphMLM-pythia8	40.99
DYJetsToLL_M-50_HT-200to400_TuneCP5_13TeV-madgraphMLM-pythia8_ext1	
DYJetsToLL_M-50_HT-400to600_TuneCP5_13TeV-madgraphMLM-pythia8	5.678
DYJetsToLL_M-50_HT-400to600_TuneCP5_13TeV-madgraphMLM-pythia8_ext1	
DYJetsToLL_M-50_HT-600to800_TuneCP5_13TeV-madgraphMLM-pythia8	1.368
DYJetsToLL_M-50_HT-800to1200_TuneCP5_13TeV-madgraphMLM-pythia8	0.6304
DYJetsToLL_M-50_HT-1200to2500_TuneCP5_13TeV-madgraphMLM-pythia8	0.1514
DYJetsToLL_M-50_HT-2500toInf_TuneCP5_13TeV-madgraphMLM-pythia8	0.003565
TTTo2L2Nu_TuneCP5_13TeV-powheg-pythia8_new_pmx	72.1
TTToSemiLeptonic_TuneCP5_13TeV-powheg-pythia8_new_pmx	314.0
TTToHadronic_TuneCP5_13TeV-powheg-pythia8_new_pmx	301.0
2018 Samples:	
ZZTo2Q2Nu_TuneCP5_13TeV_amcatnloFXFX_madspin_pythia8	6.331
WZTo2Q2Nu_TuneCP5_13TeV_amcatnloFXFX_madspin_pythia8	3.688
DYJetsToLL_M-50_HT-100to200_TuneCP5_PSweights_13TeV-madgraphMLM-pythia8	147.4
DYJetsToLL_M-50_HT-200to400_TuneCP5_PSweights_13TeV-madgraphMLM-pythia8	40.99
DYJetsToLL_M-50_HT-400to600_TuneCP5_PSweights_13TeV-madgraphMLM-pythia8	5.678
DYJetsToLL_M-50_HT-600to800_TuneCP5_PSweights_13TeV-madgraphMLM-pythia8	1.367
DYJetsToLL_M-50_HT-800to1200_TuneCP5_PSweights_13TeV-madgraphMLM-pythia8	0.6304
DYJetsToLL_M-50_HT-1200to2500_TuneCP5_PSweights_13TeV-madgraphMLM-pythia8	0.1514
DYJetsToLL_M-50_HT-2500toInf_TuneCP5_PSweights_13TeV-madgraphMLM-pythia8	0.003565
TTTo2L2Nu_TuneCP5_13TeV-powheg-pythia8	72.1
TTToSemiLeptonic_TuneCP5_13TeV-powheg-pythia8	314.0
TTToHadronic_TuneCP5_13TeV-powheg-pythia8	301.0

Table 6.3: The background MC used in the analysis with the corresponding cross section $\times$ branching fraction.

The three largest backgrounds in this analysis are

1. $t\bar{t}$, because of the presence of two real b-jets, real p_T^{miss} , and real leptons,
2. $Z + \text{jets}$, due to the real Z , jets to fake the Higgs boson, and fake p_T^{miss} , and
3. WZ and ZZ , due to the real Z , jets, and p_T^{miss} .

Each of these are modeled separately. The $t\bar{t}$ and $Z + \text{jets}$ contributions are modeled with partially data-driven methods using dedicated control regions. The background estimations are described in detail in Chapter 7.

For this analysis, we blind the Z and H mass windows in the H candidate jet soft drop mass spectrum. This allows for the analysis to be developed without bias due to the true content of the signal region. The final result is determined after the signal region is unblinded. The blinded region is divided into the signal region, $110 \text{ GeV} < m_j < 150 \text{ GeV}$ for the Higgs final state, and a $70 \text{ GeV} < m_j \leq 110 \text{ GeV}$ region, preserving the blinding for the Z boson final state.

The missing transverse momentum in the final state complicates the reconstruction of the Z' mass peak; however, observables designed for SUSY searches give a reconstruction handle. This analysis is the first use in CMS of Recursive Jigsaw Reconstruction (RJR) [4]. RJR produces mass estimators and angular variables for topologies with invisible particles in the final state. In our use case, RJR derives a mass estimator sensitive to the Z' mass, and distinct in shape from a smoothly-falling SM background.

6.3 Recursive Jigsaw Reconstruction

When the final state of a decay has particles that escape detection, the number of degrees of freedom needed to reconstruct the four-vector of the parent particle exceeds the total number of measurements available. For such events, additional constraints must be imposed on the process under study in order to reconstruct the parent particle. Recursive Jigsaw Reconstruction (RJR) is an iterative framework that applies constraints to reduce the degrees of freedom in a system [4]. The RJR rules constrain the quantities needed to boost daughter particles into the rest frames of parent particles. This reconstructs mass estimators for the parent particles while satisfying the constraints used to describe the properties of the missing particles in the final state. RJR complements established p_T^{miss} -based observables and tools, offering a separate framework for deriving quantities like $MT2$ [77], and Super Razor [78].

An illustration of the decay tree considered in this analysis is shown in Fig. 6.4 (left). The code package RestFrames implements RJR in this analysis [79]. The annotation used to describe the RJR rules follows the conventions in Ref [4].

The first two RJR rules we apply are

1. Jigsaw Rule V.1 (Invisible Rapidity)
2. Jigsaw Rule VI.1 (Invisible Mass)

The JR V.1 rule sets the rapidity of the invisible final state particles equal to the rapidity of the visible final state particles. The equivalent assumption leads to the transverse mass m_T , as used in the mass-reconstruction of a W Boson decaying to a lepton and a neutrino. The JR VI.1 rule assumes the mass of the invisible particles to be a function of the visible final state particle four vectors, with the mass being the minimum needed to satisfy all other applied Jigsaw Rules, and to ensure a non-negative mass. This allows this mass to float anywhere from 0 to equal to the mass of the visible system. These two rules drive the mass estimator of the N_D . As in the case of m_T or MT2, the distribution of N_D mass estimators from RJR, M_{ND} , produce a kinematic edge that is highly correlated with the true m_{N_D} .

The benefit of RJR is the direct interpretation between the mass estimator and the Z' mass. There are two options for the final rule used to reconstruct the Z' :

1. Jigsaw Rule VI.3 (Contra-boost Invariant)
2. Jigsaw Rule VII.3 (Invisible Minimize ΔM)

The application of Jigsaw Rule VI.3 (JR VI.3) generates a more accurate RJR Z' mass estimator ($M_{Z'}$) when the true value of m_{N_S} is large, on par with the mass of the Higgs and Z bosons. The application of Jigsaw Rule VII.3 (JR VII.3) gives a better prediction for low m_{N_S} , where $m_{N_S} \ll m_Z$. This difference comes from the assumptions the Jigsaw Rules impose. In the case of JR VI.3, the mass estimator must be invariant under anti-symmetric Lorentz boosts, instead of the normal Lorentz boost. This forces $M_{N_D} = M_{\bar{N}_D}$, which is expected under the Leptophobic Z' model. However, this removes a degree of freedom that can specify the RJR N_S mass estimator, M_{N_S} , therefore M_{N_S} tends to float to ensure the equality of JR VI.3. This pulls the mass estimator of N_S to a higher value, especially in the presence of detector resolution effects. In the case of JR VII.3, the mass difference between the two N_D particles must only be minimized, and not follow a strict equality. This defines the mass estimator for the N_S to be zero.

Figure 6.5 shows the accuracy for the JR VI.3 (Contra-Boost invariant) scenario in the $m_{N_S} = 1 \text{ GeV}$ and $m_{N_S} = 200 \text{ GeV}$ final states. Accuracy is defined as the difference between the truth value of $m_{Z'}$ and the value given by the RJR mass estimator, divided by the true mass. Figure 6.6 shows the same for the JR VII.3 reconstruction. A massive M_{N_S} , as in the case of the application of JR VI.3, more accurately reflects




Figure 6.4: Overview of recursive jigsaw reconstruction. The left image shows the decay tree and rest frames for the reconstruction analysis. Example of correlation between mass estimator and input $m_{Z'}$ for various working points (arbitrary normalization).

the truth scenario of $m_{N_S} = 200$ GeV. The JR VII.3 scenario only performs better in the $m_{N_S} = 1$ GeV samples where the mass splitting between the N_D and the N_S is small (< 500 GeV). The application of JR VII.3 forces the N_S mass estimator to zero, and more accurately reflects the truth scenario of a light, nearly massless, N_S . When the N_D and N_S mass splitting is small, m_{N_S} dictates the kinematics, so the appropriate estimation of m_{N_S} is important. When the mass splitting or N_D boost is large the final state kinematics are insensitive to m_{N_S} .

Even though the accuracy of the RJR mass estimator is better for the low mass N_S final state with JR VII.3 in some phase space, overall, JR VI.3 predicts the $m_{Z'}$ within about 10% for all mass points. This analysis focuses on boosted final state, so heavy Z' and N_D scenarios are prioritized. The contra-boost reconstruction (JR VI.3) was chosen for the RJR mass estimator, which is the final observable.

6.4 Object Specifics

Chapter 5 summarized the basics of CMS particle reconstruction, but each analysis tailors the identification and algorithms used to maximize sensitivity. The following sections describe the specific object selections and the SM boson reconstruction used in the Leptophobic Z' search.

Unless otherwise stated, data and MC treatment is the same. Common across all objects in MC is the use of *scale factors*, which account for differences between data





Figure 6.5: Accuracy of RJR Z' mass estimator ($M_{Z'}$) using the Recursive Jigsaw Reconstruction scenario JR V.1, JR VI.1, and JR VI.3, the Contra-Boost Invariant. The greener the better. The top plot is for a signal sample grid with a fixed m_{NS} of 200 GeV. The bottom is a signal sample grid with a fixed m_{NS} of 1 GeV.



parameter	abs(eta) $\leq$ 1.479	abs(eta) $>$ 1.479
$\sigma_{i\eta i\eta}$	< 0.0112	< 0.0425
$ \eta_{\text{seed}} - \eta_{\text{track}} $	< 0.00377	< 0.0067
$ \phi_{SC} - \phi_{\text{track}} $	< 0.0884	< 0.169
H/E	$0.05 + \frac{1.16}{E_{SC}} + 0.0324 \frac{\rho}{E_{SC}}$	$0.0441 + \frac{2.54}{E_{SC}} + 0.183 \frac{\rho}{E_{SC}}$
Relative Isolation	$< 0.112 + \frac{0.506}{p_T}$	$< 0.108 + \frac{0.963}{p_T}$
$ 1/E - 1/p $	< 0.193	< 0.111
Expected missing inner hits	≤ 1	≤ 1
pass conversion veto	yes	yes

Table 6.4: The criteria for the “Loose” electron ID. [61]. Subscript “SC” refers to the supercluster associated with the electron. A detailed description of the input parameters is given in section 6.4.1.

and MC reconstruction. The scale factor can be understood to be the ratio of the data efficiency over the MC efficiency for some parameter, normally binned in η and ϕ . With one notable exception described in Section 6.4.2, the scale factors applied in this analysis were centrally provided by CMS, and discussion and figures are relegated to Appendix A.

6.4.1 Electrons

The electrons used in this analysis satisfy the Loose, cut-based Electron ID summarized in Table 6.4 [61]. The ID criteria are based on parameters described in Section 5.3. The boost in the final state collimates the decay products, so no isolation requirement beyond the ones intrinsic to the Particle Flow identification are made. Electrons used in this analysis must fall within the tracker volume $|\eta| < 2.4$, and have $p_T > 20$ GeV. For about half of the 2018 data-taking period, a veto is applied on events where an electron falls in the η/ϕ region covered by the HCAL Minus Endcap sectors HEM15 and HEM16. These sectors are absent for this period, and the electron ID cannot be trusted in this region without a hadronic energy fraction measurement. The 2018 MC receives an event weight for these events to account for the lost luminosity. Chapter 8 presents an account of the impact of the loss of HEM15 and HEM16 on the reconstruction in the analysis. Scale factors are applied to account for ID [80], reconstruction [81], and trigger efficiency [82] discrepancies between MC and data.

6.4.2 Muons

Muons used in this analysis satisfy the tight muon identification (ID) [66]. The tight ID consists of Global Muon and Particle Flow muon reconstruction with additional quality cuts. The complete definition of the Tight Muon ID is summarized in Table 6.5. No isolation requirement is made to improve acceptance of muons from the boosted Z decays. Like the electrons, muons in this analysis must fall within the tracker volume $|\eta| < 2.4$, and have $p_T > 20$ GeV.

The alternative to the tight ID is the High p_T muon ID [83]. When a muon has sufficiently high momentum, it bends less in the magnetic field, and does not pass the Particle Flow criteria. However, adapting the reconstruction to focus on the outer muon system allows for the identification of these muons and a high-resolution momentum measurement. Table 6.5 also summarizes the High p_T ID definition. The main drawback of the HighPt ID is the lack of PF identification. When all other objects in the event are based on PF – like the jet-constituents and the p_T^{miss} – the lack of PF ID obfuscates the relationship between the muon and the remaining derived objects in the event. As will be seen in Section 6.4.4.1, PF ID is central to the production of the custom jet collection used in this analysis.

Tight ID muon scale factors are applied [84, 85, 86]. The ID scale factors cover the p_T range from 40 GeV to 120 GeV, and are applied as an event weight. For the dimuon channel, the event weight is the product of the two muons' scale factors. The unique combination of high- p_T triggers and tight muon ID in this analysis warranted custom muon trigger scale factors. Our scale factors further cover the signal-muon's entire kinematic p_T range from 20 GeV up to 500 GeV. Since both muons in the Z candidate could in principle fire the trigger, the applied scale factor must account for the contribution of both constituent muons to the overall trigger efficiency. To this end, the muon trigger scale factor is applied as an event weight via the following equation:

$$sf_{event} = \frac{1 - (1 - sf_{\mu 1} ef_{\mu 1})(1 - sf_{\mu 2} ef_{\mu 2})}{1 - (1 - ef_{\mu 1} ef_{\mu 2})}, \quad (6.4)$$

where $sf_{\mu i}$ is the scale factor of the constituent muon and $ef_{\mu i}$ is the corresponding MC efficiency for muons in that kinematic range. Appendix A expands on the custom trigger scale factor derivation.

6.4.3 Dimuon Z Candidate

The Z boson candidate is chosen from pairs of oppositely charged muons passing the identification criteria in Section 6.4.2, with the leading muon passing an additional p_T

parameter	Tight ID	HighPt ID
Global Muon	yes	yes
Particle Flow Muon	yes	no
χ^2 of global muon track fit	$x^2 < 10$	
Number of Muon Chamber hits in Global Muon Track	$n \geq 0$	$n \geq 1$
pT relative error of muon best track		$n < .30$
Number of Matching Muon Stations	$n \geq 1$	$n \geq 1$
Tracker Transverse Impact Parameter	$d_{xy} < 2mm$	$d_{xy} < 2mm$
Tracker Longitudinal Distance w.r.t primary vertex	$d_z < 5mm$	
Number of Pixel Hits	$n > 0$	$n > 0$
Number of Tracker Hits	$n > 5$	$n > 5$

Table 6.5: Comparison of Muon ID definitions

requirement of $p_T > 60$ GeV. If multiple pairs pass these selections, the pair closest to the Z mass and within a mass window of $70 \leq m_{ll} \leq 110$ is selected as the Z candidate.

6.4.4 Jets

This analysis requires at least one anti- k_T clustered jet with a ΔR parameter of 0.8 (AK8) with the Pile Up Per Particle Identification (PUPPI) algorithm for pile-up mitigation. While standard jets in CMS are clustered from all of the PF candidates in an event, this analysis uses a custom Particle Flow candidate collection, described in Section 6.4.4.1. To reduce the acceptance of jet-like signatures from instrumental effects, we use the tight ID for AK8PUPPI jets summarized in Table 6.6.

	2016	2017	2018
Neutral Hadron Fraction	< 0.90	< 0.90	< 0.90
Neutral EM Fraction	< 0.90	< 0.90	< 0.90
Number of Constituents	> 1	> 1	> 1
Charged Hadron Fraction	> 0	> 0	> 0
Charged Multiplicity	> 0	> 0	> 0
Charged EM Fraction	< 0.99	-	-

Table 6.6: Tight Jet ID Definitions.




Figure 6.7: The figure on the left shows the generator-level ΔR distribution at preselection for the Z and H in select signal samples. The middle and right figures show the ΔR distribution at preselection of the reco level leading (middle) or subleading (right) muon and the generator-level H. The overlap region corresponds with $\Delta R < 0.8$. Plots with reconstruction level objects were taken from samples with reclustering performed. The preselections applied are summarized in Section 6.5.

6.4.4.1 Jet Reclustering

We produce a new fat jet collection clustered from a custom PF candidate collection with the Z candidate leptons removed to preserve events where the Z and hadronic H candidates overlap. Not only does this increase the signal acceptance for select mass points and preserve isotropic angular distributions for the signal samples, it removes leptonic Z bosons faking jets.

Since high energy leptons are produced in hadron decays via the charged-weak current, standard jet collections do not clean high momentum leptons from the PF candidate collection used to cluster the jets. Without cleaning, real leptonic Z bosons fake high momentum fat jets. The most straight forward solutions are to apply isolation requirements to the lepton objects, veto jets with high lepton content, or to require separation between the Z candidate and the jet. While these options are effective for most topologies, for final states where there is true overlap between the leptonic Z and the hadronic H, signal would be lost. Figure 6.7 shows the ΔR spectra between the Z candidate and the generator Higgs Boson for select signal points.

To retain signal in the presence of overlap, we remove the PF candidates that build the Z candidate from the standard PF candidate collection. A new fat jet collection is then created from this reduced group using the anti- k_T algorithm with PUPPI for a jet radius of 0.8. This creates a new AK8PUPPI jet collection with all jet-properties like soft drop computed.

Figure 6.8 shows the kinematic effects of the reclustering on jets in the overlap region. The reclustering changes the shapes of the p_T and η distributions, revealing the skew

in the jet kinematics introduced by the leptons originating from the Z candidate decay. The ΔR spectrum between the reconstructed Z and H candidates at preselection in the signal region is shown in Fig. 6.9. From the comparison of the custom jet collection and the standard collection distributions in the aforementioned figures, we can clearly see the reduced background acceptance due to the custom collection. Reclustering removes the Z faking the H in the signal final state.




Figure 6.8: Kinematic plots of the $m_{Z'} 3000 \text{ GeV}$ $m_{N_D} 1200 \text{ GeV}$ $m_{N_S} 1 \text{ GeV}$ signal mass point at skim level selections, in the Z candidate and H candidate overlap region ($\Delta R(Z, H) < 0.8$). Skims selections require the Z candidate reconstruction, $p_T^Z > 150.0 \text{ GeV}$, and at least one fat jet. H candidate chosen as the fat jet closest to the Higg's mass. Distributions are normalized to unit area to show shape differences.

6.4.4.2 Jet Tagging

Jet substructure analysis allows for flavor identification of jets. For example, jets originating from b-mesons (b-jets) have displaced vertices due to the longer lifetimes of b-quarks. Finding these secondary vertices identifies a jet as a b-jet. Reconstructing and making selections on structures like secondary vertices have long been a feature of collider physics analysis.




Figure 6.9: ΔR distribution between the Z candidate and H candidate at preselection in the H channel signal region for a few benchmark signal mass distributions. Scaled to full Run 2 luminosity, and a cross section of 10 fb^{-1} . Left figure is reclustered. Right figure is not reclustered.

Taggers based on machine learning (ML) can correlate the different jet-substructure variables to produce more nuanced identification and can also take lower-level observables –like the PF candidates themselves– to increase the tagger-yield. In this way, taggers targeting specific decay modes and particles can be developed.

Numerous flavor-taggers using ML exist in CMS to enhance the identification and selection of hadronically decaying objects [87]. The ML-based taggers in the boosted-regime outstrip their cut-based counterparts, and have become the default taggers in CMS analyses.

For this analysis, we seek to enhance the selection of fat jets originating from Higgs Boson decays to two b-quarks. Figure 6.10 compares the performance of different CMS $H \rightarrow b\bar{b}$ taggers in this analysis. All four of the analysis-appropriate taggers are based on Deep Neural Networks (DNNs). The DeepMassDecorrelHbbvQCD and DeepMassDecorrelZHbbvQCD are by far the most performant of the available taggers. For signal efficiencies of about 80%, both the DeepMassDecorrelTagZHbbvsQCD and the DeepMassDecorrelTagHbbvsQCD offer similar performance. To facilitate the addition of future hadronic Z channels, the DeepMassDecorrelTagZHbbvsQCD tagger is used.

This class of deep-taggers, colloquially called “DeepAK8” in CMS, are a class of taggers based on low-level event content [87]. The tagger takes forty-two attributes (p_T , charge, track quality, etc.) of each of the 100 highest momentum PF candidates in the event and combines this with information from up to seven secondary vertices. The secondary vertex information especially aides heavy-flavor discrimination. The DeepAK8 algorithms are designed for fat jets with $p_T > 200 \text{ GeV}$. We use the




Figure 6.10: Comparison of jet tagger efficiency for double-b or Higgs taggers. Background efficiency is defined as the weighted sum of background events passing all selections and b-tagging over the weighted sum of background events passing selections alone. Signal efficiency is defined the same way. The signal sample used is a privately generated $m_{Z'} = 2000$ GeV, $m_{N_D} = 500$ GeV, $m_{N_S} = 200$ GeV sample. The selections used were $p_T(Z) > 200$ GeV, $p_T(H) > 300$ GeV, and $p_T^{miss} > 200$ GeV.

mass-decorrelated (MD) version of the tagger, with the decorrelation step done via adversarial training. Without mass decorrelation, the jet mass sidebands would be less useful as control regions for the analysis.

The DeepAK8MDZHbb score is defined as

$$score_{\text{DeepAK8MDZHbb}} = \frac{P_{H \rightarrow b\bar{b}} + P_{Z \rightarrow b\bar{b}}}{P_{H \rightarrow b\bar{b}} + P_{Z \rightarrow b\bar{b}} + P_{H \rightarrow c\bar{c}} + P_{Z \rightarrow c\bar{c}} + P_{QCD}}, \quad (6.5)$$

where each P value is the prediction score assigned by the algorithm to each process [88]. The loose working point of 0.8 is used in this analysis. Scale factors for b-tagging are applied, and exist up to a jet- p_T of 700 GeV.

6.4.5 Higgs Candidate Reconstruction

From the custom collection of fat jets, the Higgs candidate is chosen as the jet with soft drop mass closest to Higgs mass and a DeepMassDecorrelTagZHbbvsQCD score > 0.8 . To further eliminate jets not originating from boson decays, we require a soft drop mass cut of $m_j > 30$ GeV. The soft drop mass is used to define our signal and sideband regions; therefore, the Higgs Boson mass window, the final selection, $110 \text{ GeV} < m_{hcand} < 150 \text{ GeV}$, is blinded.

6.4.6 Missing Transverse Momentum

In this analysis, we use type-1 corrected PF $\vec{p}_T^{\text{miss}}$ (as described in Section 5.5). There is an artificial and undesired sinusoidal fluctuation in the reconstructed $\vec{p}_T^{\text{miss}}$ in CMS [89]. The period of the oscillation is roughly 2π , and is possibly due to offline detector regions, detector misalignment, or the displacement of the beam spot. Our analysis uses the ϕ coordinate of $\vec{p}_T^{\text{miss}}$ in RJR, so the $\vec{p}_T^{\text{miss}}$ is corrected to avoid artificial bias [89].

6.5 Preselection Cut Scheme

The preselection cut scheme can be summarized as follows:

1. At least one fat jet with $|\eta| < 2.4$, DeepAK8ZHbbMD > 0.8 , $p_T > 300.0$, and $m_{sd} > 30$
2. At least one oppositely charged dilepton pair (each lepton with $|\eta| < 2.4$) within the Z mass window $70 \text{ GeV} < m_j \leq 110 \text{ GeV}$ with $p_T > 100.0 \text{ GeV}$.
3. $p_T^{\text{miss}} > 75.0$

These cuts were chosen to fully capture the $H \rightarrow b\bar{b}$ decay in a jet of radius 0.8, to preserve the background estimation, and to maximize the expected statistical sensitivity of the analysis. The radius of a fat jet with two subjets (as in the case of a fully-merged $H \rightarrow b\bar{b}$) goes roughly as

$$\Delta R \approx \frac{2m_{\text{parent}}}{p_T^{\text{parent}}}. \quad (6.6)$$

For the SM Higgs, the two daughter b-jets are merged in a jet of radius 0.8 when the parent H candidate has a p_T of roughly 300 GeV. This analysis uses soft drop mass

sidebands, so it is paramount to capture the full Higgs decay in the jet radius to have signal-free sidebands. Thus, we settle on a H candidate $p_T > 300$ GeV. The partially data-driven estimation for $Z + \text{jets}$ described in detail in Section 7.2 requires fitting a background model to the data sideband and MC signal region. To maintain the quality of the fits, the signal region requires around 200 background events. Figure 6.11 shows the signal and sideband MC yields for the potential grid of p_T^{miss} and Z candidate p_T cuts, keeping the H candidate p_T cut greater than 300 GeV. Cut-values outside of the grid are either too-tight a priori, or are unnecessarily loose. Within the grid, expected limits with minimal systematics to gauge sensitivity were tested for a sample of signal points. The lowest limit across the bulk of the signal points between two cut-schemes was chosen to be compared to the next cut-scheme, until the cut scheme with the greatest sensitivity was found. Consider the following two scenarios:

1. Z candidate $p_T > 100$ GeV, $p_T^{\text{miss}} > 75$ GeV
2. Z candidate $p_T > 100$ GeV, $p_T^{\text{miss}} > 50$ GeV

The differences between the two expected limits is shown in Fig. 6.12. The first cut scheme gives the best results over the entire cut-scheme grid, for most of the signal points considered. Intuitively, this performance is expected. The selected cut scheme has the average background yield of the considered schemes. The considered schemes were chosen to cover a window around the minimum number of background events needed to preserve the closure of the background estimation method. Thus, the scheme with the average yield corresponds to the point where the background is minimized without diminishing the quality of the background estimation. Table 6.7 shows the signal yields for four benchmark signal samples with this cut-scheme for a signal cross section normalized to 10 fb.

Figure 6.13 shows selected kinematic distributions of the final state objects in the soft drop mass sidebands ($30 \text{ GeV} < m_{\text{higgscand}} < 70 \text{ GeV}$ and $150 \text{ GeV} \leq m_{\text{higgscand}}$) at preselection level. Figure 6.14 shows the RJR mass estimator distributions in the sidebands. The plots show the comparison between data and MC. Figure 6.15 shows select signal region particle kinematic spectra, and Fig. 6.16 shows the signal region Z' , N_D , and N_S RJR mass estimators.




Figure 6.11: The cutflow of total background in the sideband and the signal region for different Z candidate p_T and p_T^{miss} cut values. The H candidate has a $p_T > 300$ GeV. The grid was chosen to maintain the Z + jets and $t\bar{t}$ yields necessary to maintain the Z + jets estimation method. Roughly 200 background events are needed in the signal region.




Figure 6.12: Percent difference between expected limits for the cut scenarios Z candidate $p_T > 100$ GeV, $p_T^{\text{miss}} > 75$ GeV and Z candidate $p_T > 100$ GeV, $p_T^{\text{miss}} > 50$ GeV. Negative values indicate that the limit of the $p_T^{\text{miss}} > 75$ GeV scenario is lower (more sensitive) than the limit of the $p_T^{\text{miss}} > 50$ GeV.

cut description	mzp1500 mnd600mns200	mzp2000 mnd800mns200	mzp3000 mnd400mns200
Expected Events*	1012.7	1012.7	1012.7
Events Passing Trigger*	943.3	963.3	946.8
Events with $70 < m_{\ell\ell} < 110^*$	763.42	793.9	708.4
Events with at least 1 fat jet*	474.95	637.53	620.37
Events passing H reco*	237.20	412.51	445.8
Events Passing $MET > 75.0$	216.94	392.25	405.56
Events Passing $Zp_T > 100.0$	206.57	384.85	393.46
Events Passing Fat Jet $p_T > 300.0$	172.23	351.72	372.03
Events Passing 0.8 btag WP	127.63	276.13	299.78
Events in Signal Region	91.44	191.29	204.33

Table 6.7: Signal yields with luminosity scaling for 2017 and 2018. The cross section is normalized to 10 fb. Cuts with an asterisk do not have btagging or muon ID scale factors applied.




Figure 6.13: The soft drop mass sideband region plots of select final state particle object kinematic distributions. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.




Figure 6.14: The soft drop mass sideband plots at preselection featuring the RJR mass estimator spectra. The plots are in units of GeV. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.




Figure 6.15: The soft drop mass signal region particle object p_T spectra. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.




Figure 6.16: The soft drop mass signal region plots at preselection featuring the RJR mass estimator spectra. Preselections are described in Section 6.5. The plots are in units of GeV. The normalization for Drell-Yan MC derived in Section 7.2.1 is applied, but otherwise the filled histograms are nominal simulation.

Chapter 7

Background Estimations

The three main backgrounds for this analysis are $t\bar{t}$ production, $Z + \text{jets}$ events from the Drell-Yan process, and diboson production. While mitigated as much as possible through the event selection described in Chapter 6, we must extrapolate the remaining background contribution to the RJR Z' mass estimator in the signal region. The background can be parameterized by the contributions of the individual sources,

$$N_{SR}(M_{Z'}) = N_{SR}^{t\bar{t}} + N_{SR}^{Z+\text{jets}} + N_{SR}^{VV},$$

where the $N_{SR}^{t\bar{t}}$, $N_{SR}^{Z+\text{jets}}$, and N_{SR}^{VV} are detailed in Sections 7.1, 7.2, and 7.3, respectively.

This analysis prioritizes the use of data-driven background estimation techniques. Both the $t\bar{t}$ and $Z + \text{jets}$ estimations use partially data-driven methods. Data-driven methods reduce the impact of systematic uncertainties that arise from simulation and inform the background estimations when simulation statistics are low or the models are poorly constrained.

7.1 $t\bar{t}$ Estimation

The leptonic and semileptonic decay modes of SM $t\bar{t}$ production, illustrated in Fig. 7.1, are the main backgrounds in this analysis. When both of the W bosons from the top quark decay to a muon and a neutrino, the $t\bar{t}$ final state matches the Z' to anomalon final state: two b -jets, real p_T^{miss} , and a real dimuon. Semileptonic $t\bar{t}$ contributes because the custom jet collection allows muons to be within the jet radius. This can allow high-momentum muons from the decay of the b -quark via a virtual W boson to contribute a second muon to mimic a Z candidate. In this case, the hadronic top quark fakes the Higgs candidate.

We use a data-driven background estimation adapted from the leptonic $t\bar{t}$ estimation in [90]. Same lepton-flavor final states comprise only half of the leptonic $t\bar{t}$ decay width. Leptonic $t\bar{t}$ decays to mixed electron and muon ($e\mu$) final states as well, and these can serve as a proxy for the same flavor final-states. The $e\mu$ channel is similar




Figure 7.1: Illustrative Feynman diagrams of leptonic (left) and semileptonic (right) $t\bar{t}$ decay channels. In the leptonic-case, the two leptons from the leptonic W decay fake the Z candidate and the two b -jets merge to fake the Higgs candidate jet. In the semileptonic case, a lepton from the leptonic W fakes one of the leptons in the Z candidate, but the second lepton comes from the decay of the b -quark via a virtual W boson. The hadronic top produces the fat jet candidate.

in cross section and kinematics, and therefore offers an orthogonal decay channel to extract the $t\bar{t}$ contribution in the signal region.

To reconstruct the $e\mu$ channel of the $t\bar{t}$ background, each event is scanned for both Z ($\ell\ell$) and Z ($e\mu$) candidates, with accepted events only containing one type of Z candidate. The kinematic requirements in the Z reconstruction are the same across both the $\mu\mu$ and $e\mu$ channels. The electron selections, muon selections, and Z candidate reconstruction are described in Section 6.4. Either the electron or muon can be the leading or subleading lepton. All leptons used to reconstruct Z candidates ($\ell\ell$ and $e\mu$) are removed from the PF candidate collection used to build the custom jet collection described in Section 6.4.4.1. Since high p_T leptons are rare, it is unlikely that an event has two types of Z candidate, so few events are lost by reconstructing both lepton pairs in an event. See Table 7.1 for a list of event yields with more than one type of Z candidate.

To build the $e\mu$ channel both single electron and single muon triggers are used. Table 6.2 summarizes the triggers considered. Since CMS data is divided into sets based on triggers, both the SingleElectron/EGamma and SingleMuon data sets are used to build the data control region to extract the $e\mu$ shape. Once the data sets are processed and events have passed the corresponding data sets' respective triggers, the resulting yields are combined, and any duplicates between the two data sets (events that fired both an electron and a muon trigger) are dropped. Unlike data, simulation contains

Z candidates present	Event yields without trigger	Event yields with trigger
Z($\mu\mu$) only	166035	159406
Z($\mu\mu$) and Z(ee)	4	4
Z($\mu\mu$) and Z($e\mu$)	15	4
Z($\mu\mu$) and Z(ee) and Z($e\mu$)	1	1
Z(ee) only	654	133
Z(ee) and Z($e\mu$)	21	0
Z($e\mu$) only	9903	3581

Table 7.1: Reconstructed Z candidate yields with > 0 fat jets and after p_T^{miss} filter application in the 2018D SingleMuon data. Z ($e\mu$) and Z ($\ell\ell$) candidates are reconstructed for each event. This table shows the number of events with multiple Z candidates reconstructed, which are vetoed. The overlap region is negligible.

all trigger paths, and the trigger requirement is simply an “or” of the electron and muon triggers for a given year. The trigger scale factor is applied based on the leading lepton of the flavor that fired the trigger. If both flavor triggers are fired, the trigger scale factor is applied based on the muon in the event, since it is serving as the background estimation for the dimuon channel.

Once the $e\mu$ “Z” candidate is built, the selections proceed as in the dimuon channel. The $e\mu$ candidate events in data that pass the preselection cuts and the fall within the soft drop mass signal region provide the shape of the RJR Z’ mass estimator distribution for the $t\bar{t}$ background estimation. The $e\mu$ shape in the soft drop mass signal region is normalized to the expected $\mu\mu$ channel yields via the ratio of $\mu\mu$ to $e\mu$ yields in simulation. The $\mu\mu$ to $e\mu$ ratio is described in more detail in Section 7.1.1.

7.1.1 $e\mu$ channel normalization

We assume that the $e\mu$ channel only differs in yield from the standard $\mu\mu$ channel. There are two ways to calculate the normalization scale factor:

1. Use a $t\bar{t}$ control region in data to calculate the normalization.
2. Use the ratio of the $\mu\mu$ and $e\mu$ $t\bar{t}$ simulation yields to generate the normalization.

This analysis uses the second option, where the scale factor is derived as the ratio of the $\mu\mu$ and $e\mu$ $t\bar{t}$ simulation yields. The simulation derivation offers a large statistical




Figure 7.2: Value of $e\mu/\mu\mu$ in simulation versus cut for leptonic and semileptonic $t\bar{t}$. Overall, the $e\mu$ and $\mu\mu$ channels of $t\bar{t}$ respond similarly to the selections, producing a flat cutflow. This changes with the DeepAK8 tagger requirement in the leptonic channel. This produces about a 20% difference in efficiency between the $e\mu$ and $\mu\mu$ channels.

sample in an otherwise statistically limited region of phase space. To have a sample in data with a statistical uncertainty comparable to the sample in simulation, the b-tagging requirement would have to be released. The b-tagging primarily removes non-resonant WW production, therefore releasing the b-tagging requirement contaminates the $t\bar{t}$ control region in data. The impact of the additional uncertainties associated with simulation is reduced by cancellation in the $\mu\mu$ to $e\mu$ ratio.

The $\mu\mu$ to $e\mu$ ratio is derived from the full soft drop mass spectrum at preselection in simulation. While looser selections would reduce the statistical uncertainty on the ratio, a roughly 20% difference in the DeepAK8 tagging efficiency between $\mu\mu$ and $e\mu$ final state leptonic $t\bar{t}$ demands the entire cut-scheme to capture the yield differences due to the tagger. Figure 7.2 shows the progression of the $e\mu$ over $\mu\mu$ ratio through the selections for leptonic and semileptonic $t\bar{t}$ simulation. The DeepAK8 tagger is agnostic to the difference in final state for semileptonic $t\bar{t}$ because the fat jet candidate is from the hadronic top. The hadronic top is back-to-back with the leptons, so the leptons have no physical proximity to the jet. This is not always the case with the leptonic $t\bar{t}$. Figure 7.3 shows the spatial difference in the pseudorapidity-azimuth plane between the selected fat jets and the dilepton candidates in the $e\mu$ $t\bar{t}$. The selected muons and electrons are removed from the jet clustering, but the response of the algorithm indicates that electron and muon jet constituents have different impacts when generating the DeepAK8 score.




Figure 7.3: Distribution of ΔR between the leading lepton in the $e\mu$ fake “Z” candidate and the selected Higgs boson candidate fat jet. Semileptonic $t\bar{t}$ is back-to-back, while leptonic $t\bar{t}$ has events where the dilepton and the fat jet overlap.

If background contributions are from leptonic $t\bar{t}$ alone, the expected $\mu\mu$ to $e\mu$ ratio is 1:2. This ratio assumes that tight muon selection and loose electron selection are equally efficient; however, the real ratio of tight muon to loose electron selection in this analysis is roughly 1:0.7. With these efficiencies, the anticipated leptonic ratio becomes 1:1.4. When one considers the contributions of semileptonic $t\bar{t}$, the ratio for the total derivation region becomes

$$\frac{\mu\mu_{\text{total}}}{e\mu_{\text{total}}} = \frac{\mu\mu_{\text{semilep}} + \mu\mu_{\text{lep}}}{e\mu_{\text{semilep}} + e\mu_{\text{lep}}} . \quad (7.1)$$

To arrive at the expected ratio, we rearrange Eq. 7.1 in terms of known quantities,

$$\begin{aligned} \frac{\mu\mu_{\text{total}}}{e\mu_{\text{total}}} &= \frac{\mu\mu_{\text{semilep}}}{e\mu_{\text{semilep}} + e\mu_{\text{lep}}} + \frac{\mu\mu_{\text{lep}}}{e\mu_{\text{semilep}} + e\mu_{\text{lep}}} , \\ &= \frac{1}{\frac{e\mu_{\text{semilep}}}{\mu\mu_{\text{semilep}}} + \frac{e\mu_{\text{lep}}}{\mu\mu_{\text{semilep}}}} + \frac{1}{\frac{e\mu_{\text{semilep}}}{\mu\mu_{\text{lep}}} + \frac{e\mu_{\text{lep}}}{\mu\mu_{\text{lep}}}} . \end{aligned} \quad (7.2)$$

The result of Eq. 7.2 can be simplified with $\mu\mu_{\text{semilep}}/e\mu_{\text{semilep}} = 0.75$, taken from simulation, and $\mu\mu_{\text{lep}}/e\mu_{\text{lep}} = 1.4$, which comes from the combinatorics and efficiencies explained above. With these substitutions, Eq. 7.2 becomes

$$\frac{\mu\mu}{e\mu} = \frac{1}{0.75 + \frac{e\mu_{\text{lep}}}{\mu\mu_{\text{semilep}}}} + \frac{1}{\frac{e\mu_{\text{semilep}}}{\mu\mu_{\text{lep}}} + 1.4}. \quad (7.3)$$

From studies with simulation, we also know

- $\mu\mu_{\text{semilep}} = 0.66 \mu\mu_{\text{total}}$,
- $\mu\mu_{\text{lep}} = 0.34 \mu\mu_{\text{total}}$,
- $e\mu_{\text{semilep}} = 0.6 e\mu_{\text{total}}$, and
- $e\mu_{\text{lep}} = 0.4 e\mu_{\text{total}}$.

With these substitutions, Eq. 7.3 becomes

$$\frac{\mu\mu_{\text{total}}}{e\mu_{\text{total}}} = \frac{1}{0.75 + \frac{0.4 e\mu_{\text{total}}}{0.66 \mu\mu_{\text{total}}}} + \frac{1}{\frac{0.6 e\mu_{\text{total}}}{0.34 \mu\mu_{\text{total}}} + 1.4}. \quad (7.4)$$

Now everything is in terms of $\mu\mu_{\text{total}}$ and $e\mu_{\text{total}}$. Multiplying both sides of Eq. 7.4 by $e\mu_{\text{total}}/\mu\mu_{\text{total}}$ gives

$$1 = \frac{1}{0.75 \frac{\mu\mu_{\text{total}}}{e\mu_{\text{total}}} + 0.6} + \frac{1}{1.8 + 1.4 \frac{\mu\mu_{\text{total}}}{e\mu_{\text{total}}}}. \quad (7.5)$$

We have now isolated $\mu\mu_{\text{total}}/e\mu_{\text{total}}$. Setting $x = \mu\mu_{\text{total}}/e\mu_{\text{total}}$ yields,

$$1 = \frac{1}{0.75x + 0.6} + \frac{1}{1.8 + 1.4x}, \quad (7.6)$$

which can be rearranged to give

$$\begin{aligned} 1.8 + 1.4x &= \frac{1.8 + 1.4x}{0.75x + 0.6} + 1, \\ (0.75x + 0.6)(1.8 + 1.4x) &= 1.8 + 1.4x + 0.75x + 0.6, \\ 1.05x^2 + 2.19x + 1.08 &= 2.15x + 2.4, \\ 1.05x^2 + 0.04x - 1.32 &= 0. \end{aligned} \quad (7.7)$$

Taking the physical solution of Eq. 7.7, we receive an expected ratio of

$$x = \frac{\mu\mu_{\text{total}}}{e\mu_{\text{total}}} = 1.1. \quad (7.8)$$




Figure 7.4: A comparison of the shapes of the $t\bar{t}$ MC in the $e\mu$ and $\mu\mu$ channels, with the normalization scale between them. This yield is expected based on the relative inefficiency between muon and electron identification, and the make-up of our control region.

Figure 7.4 shows comparisons of basic kinematic distributions in the $\mu\mu$ and $e\mu$ $t\bar{t}$ simulation, and the derived $\mu\mu$ to $e\mu$ ratio (taken directly from the simulation yields) of $\mu\mu/e\mu = 1.094 \pm 0.03$ (statistical uncertainty). The Eq. 7.8 result is in reasonable agreement with the measured ratio.

7.1.2 $e\mu$ -channel informed estimation of $\mu\mu$ background

The final step is to apply the normalization to the $e\mu$ channel distributions. Figure 7.5 shows the agreement between the full-Run 2 $e\mu$ estimation and the Run 2 $t\bar{t}$ $\mu\mu$ simulation in two preselection observables. The RJR Z' mass estimator distribution in Fig. 7.5 (right) represents the $t\bar{t}$ estimation in the signal region. The agreement between the estimation and the signal region $\mu\mu$ simulation is good, and verifies the performance of the method. Even with the low data-event yield, using the data-driven method instead of a simulation-only estimation dramatically reduces the impact of systematic uncertainties, despite the increase in statistical uncertainties. The uncertainties associated with the $t\bar{t}$ background estimation are discussed in more detail in chapter 8.

Additional control plots for the $e\mu$ channel are included in Appendix B.




Figure 7.5: Overlays of the normalized $e\mu$ extrapolation (data set with scaling) and the $\mu\mu$ channel MC for preselection cuts. The plot on the left is the soft drop mass distribution. The right plot is the $t\bar{t}$ background estimation for the recursive jigsaw mass estimator of the Z' in the soft drop mass signal region.

7.2 α -Method for $Z + \text{jets}$ Estimation

Drell-Yan $Z + \text{jets}$ production is the second-largest background for this analysis. To estimate the $Z + \text{jets}$ contribution in the signal region RJR Z' mass estimator ($M_{Z'}$) spectrum, we build a transfer function called the α ratio, $\alpha(M_{Z'})$, which is the ratio of the sideband (SB) $Z + \text{jets}$ shape $N_{SB}^{MC,Z+jets}(M_{Z'})$ in MC to the signal region (SR) simulation shape $N_{SR}^{MC,Z+jets}(M_{Z'})$,

$$\alpha(M_{Z'}) = \frac{N_{SR}^{MC,Z+jets}(M_{Z'})}{N_{SB}^{MC,Z+jets}(M_{Z'})}. \quad (7.9)$$

After the other background contributions are subtracted from the data sidebands, the α ratio is applied to the data sidebands to correct for the differences between sideband and signal region shape and yield. This procedure is known as the α -method, and produces the $Z + \text{jets}$ estimation in the signal region.

The α -method has four steps:

1. Normalize the background $Z + \text{jets}$ simulation using data sidebands.
2. Fit the $M_{Z'}$ $Z + \text{jets}$ MC signal and sideband region spectrums with smoothly falling functions to derive the α ratio in Eq. 7.9,

3. Fit the sideband data $M_{Z'}$ spectrum with a smoothly falling function and subtract the $t\bar{t}$ and VV contributions to estimate the $Z + \text{jets}$ contribution in data,

$$N_{SB}^{data,Z+jets}(M_{Z'}) = N_{SB}^{data}(M_{Z'}) - N_{SB}^{data,t\bar{t}}(M_{Z'}) - N_{SB}^{data,VV}(M_{Z'}).$$

4. Multiply of the data sideband distribution, corrected to represent only the $Z + \text{jets}$ background, of the $M_{Z'}$ spectrum $N_{SB}^{data}(M_{Z'})$ by the α ratio for a data-dependent estimation of the $Z + \text{jets}$ contribution in the signal region, $N_{SR}^{est,Z+jets}(M_{Z'})$,

$$N_{SR}^{est,Z+jets}(M_{Z'}) = N_{SB}^{data,Z+jets}(M_{Z'})\alpha(M_{Z'}).$$

Section 7.2.1 and 7.2.2 cover the steps in more detail.

The regions in the p_T^{miss} -soft drop mass plane used in the α -method are shown in Fig. 7.6. The signal region is as described in Section 6.5: a fat jet with $p_T > 300.0$ GeV, a DeepAK8 working point > 0.8 and within the Higgs boson mass window; the Z candidate $p_T > 100.0$ GeV; the $p_T^{\text{miss}} > 75.0$ GeV; and an $|\eta| < 2.4$. The sidebands used for the normalization of the Drell-Yan simulation and the sideband distribution in the α ratio itself are the Higgs candidate soft drop mass sidebands but without a p_T^{miss} requirement. This opens the soft drop mass signal region with an inverted p_T^{miss} requirement as a validation region for the method.

7.2.1 Background Normalization

Before deriving the α ratio, the Drell-Yan $Z + \text{jets}$ simulation is normalized to the data sideband $Z + \text{jets}$ yield. $Z + \text{jets}$ is characterized by low amounts of p_T^{miss} that originate from jet mismeasurement, so the p_T^{miss} cut is released to enrich the soft drop mass sidebands with $Z + \text{jets}$. The normalization is derived by fitting the soft drop mass spectrum. The functional shape is taken as the sum of templates of the individual backgrounds in MC, with an additional parameter A for the $Z + \text{jets}$ normalization,

$$f_{\text{total}}(m_{\text{softdrop}}) = Af_{Z+jets}(m_{\text{softdrop}}) + f_{t\bar{t}}(m_{\text{softdrop}}) + f_{VV}(m_{\text{softdrop}}). \quad (7.10)$$

The shapes used to determine the background templates are summarized in Table 7.2. The WZ and ZZ samples are fit as one VV shape, $f_{VV}(m_{\text{softdrop}})$. The VV spectrum is fit with a Gaussian peak on top of a linearly falling background to capture the Z peak and the W contributions, respectively. The $t\bar{t}$ ($f_{t\bar{t}}(m_{\text{softdrop}})$) is modeled as the sum of two Gaussians: a narrow peak centered at the top mass, and a broad




Figure 7.6: Illustration of regions used in the α -method. All considered regions have a fat jet with $p_T > 300.0$ GeV and a DeepAK8 working point > 0.8 , the Z candidate $p_T > 100.0$ GeV, and an $|\eta| < 2.4$. The SR label indicates the signal region, SB the sideband, VR the validation region, and ZR the blinded region around Z peak.

Gaussian centered at lower mass for the top events that are not-fully reconstructed in the AK8 jet cone. Finally, the Z + jets shape is a fifth degree polynomial. Figure 7.7 shows the individual background fits. Equation 7.10 is then fit to the data sidebands with A allowed to float and the template parameters fixed. Figure 7.8 shows the normalization and the sideband fit.

7.2.2 RJR Background Shape Parameterization and Estimation

After the Z + jets MC is normalized, the next steps are to build the transfer function and extrapolate to the signal region. The α -method requires the following background MC fits:

1. Drell-Yan MC RJR Z' mass estimator distribution in the SR. This fit defines the numerator in the α ratio.
2. Drell-Yan MC RJR Z' mass estimator distribution in the SB. This fit defines the denominator in the α ratio.
3. $t\bar{t}$ distribution in the RJR Z' mass estimator in the sideband region. This fit is subtracted from the data sideband distribution to isolate the Drell-Yan

Background	Shape Function
Z + jets	$f_{Z+jets}(x) = a + bx + cx^2 + dx^3 + gx^4 + hx^5$
$t\bar{t}$	$f_{t\bar{t}}(x) = A(De^{-0.5(\frac{x-b}{c})^2} + (1-D)e^{-0.5(\frac{x-g}{h})^2})$
VV	$f_{VV}(x) = A(e^{-0.5(\frac{x-b}{c})^2} + Dx + G)$
total sideband	$f_{total}(x) = Af_{Z+jets} + f_{t\bar{t}} + f_{VV}$

Table 7.2: Background shapes for Z + jets MC simulation normalization. The three background templates' ($f_{Z+jets}(x)$, $f_{t\bar{t}}(x)$, and $f_{VV}(x)$) parameters are fixed after fits to the respective MC. The total function $f_{total}(x)$ with the fixed background parameters is fit to the data sidebands with the normalization A allowed to float.




Figure 7.7: Individual backgrounds in the soft drop mass spectrum, and their respective fits. Distributions shown at preselection with the p_T^{miss} cut released, with no SR and SB separation.




Figure 7.8: Stacked plots showing the full soft drop mass sideband fit. Left stack has the $Z + \text{jets}$ plotted without the derived normalization, the right stack has the normalization applied. Both plots include the full background template fit to MC only, and the data sideband fit.

contribution in data.

4. Diboson distribution in the RJR Z' mass estimator in the sideband region. This fit is subtracted from the data sideband distribution to isolate the Drell-Yan contribution in data.

Lastly, there is a fit to the data SB, from which the $t\bar{t}$ and VV are subtracted. A simple two-parameter exponential fit function is used. Figure 7.9 shows the background fits and the α ratio. The fitting region begins at the $M_{Z'}$ value corresponding to the distribution maximum, to avoid the Drell-Yan + jets turn-on in the mass estimator spectrum. The fitting region ends when the bin yield is ≤ 0.1 event.

Figure 7.10 shows the agreement between the individual background fits and sideband data, and the data sideband fit. Finally, we subtract the $t\bar{t}$ and VV shapes from the fit to the sideband data. The subtracted distribution is shown in Fig. 7.11. The functional difference is multiplied by the α ratio, and this product is taken as the $Z + \text{jets}$ contribution in the signal region. The resulting estimation is shown in Fig. 7.12. The systematics associated with the α -method are described in Chapter 8.




Figure 7.9: Distributions used in α -method background estimation. The lower range of the fitted region implies a final cut at 1400. The red bands are the fit uncertainties. Top Left: Sideband Drell-Yan + jets MC RJR Z' mass estimator distribution. Top Middle: RJR mass estimator sideband of $t\bar{t}$. Top Right: WZ and ZZ sideband RJR mass estimator. Bottom Left: Signal region Drell-Yan + jets MC RJR mass estimator distribution. Bottom Right: α ratio formed from the ratio of the fits in the bottom left and top left plots. Bottom Middle: Data sideband and fit.




Figure 7.10: Plots showing the fit agreement with the data sideband. The left plot shows the sideband data overlaid with the individual background fits. The right plot shows the agreement between the sideband data and the fit.




Figure 7.11: The subtracted data sideband distribution. The black line represents the $Z + \text{jets}$ contribution in the data sideband. This plot is the difference between the data sideband fit, and the $t\bar{t}$ and VV fits. Uncertainty bands are discussed in Chapter 8.




Figure 7.12: The signal region $Z + \text{jets}$ estimation. Left plot shows the estimation on a linear scale in the RJR mass estimator region between 1400 and 3000 GeV. Sideband data events exist in this region. The left plot shows the entire extrapolation range. Drell-Yan MC is included to guide the eye, and is not expected to agree. At high p_T , we rely on data, rather than simulation, to inform the shape. Uncertainty bands are discussed in Chapter 8.




Figure 7.13: Top Left: Sideband Drell-Yan + jets MC RJR Z' mass estimator distribution with derived normalization. Top Middle: Sideband of $t\bar{t}$. Top Right: WZ and ZZ sideband. Bottom Left: validation region Z + jets MC. Bottom Right: α ratio formed from the ratio of the fits in the bottom left and top left plots. Bottom Middle: data sideband and fit.

7.2.3 α -method validation

The validation region is used to test the performance of the α -method before it is applied to the signal region. The validation region is unblinded; therefore, data can be compared with the extrapolation. The same normalization is used in the validation as the nominal background estimation. Figures 7.13 through 7.14 show the sideband and validation region α ratio shapes and results. The α -method provides an estimation of the Z + jets appropriate to the validation region data in our area of interest. Validation region data and MC agreement plots are included in Appendix C.

7.3 Diboson Estimation

The diboson (VV) background is the final background in this analysis and is taken directly from MC. The $b\bar{b}$ requirement excludes WW production, and leaves only the semileptonic final states of WZ and ZZ production. Collectively, these processes contribute about two percent of the background in the signal region. With such low yields, we neglect VH production due to its lower cross section.




Figure 7.14: Plots showing the application of the α -method in the validation region. The left plot shows the fit to the sideband data with the $t\bar{t}$ and VV components subtracted. The middle plot is the α ratio between the sidebands and the validation region. The right plot is the estimation of the $Z + \text{jets}$ contribution in the validation region through the multiplication of the functions in the two plots to the left. MC included to guide the eye. Perfect agreement between data and MC is not expected in this data-driven background.

Chapter 8

Systematic Uncertainties

This section accounts for the systematic uncertainties present in this analysis. For each systematic, the entire analysis flow is repeated with the respective up or down variation applied. Depending on the systematic, this can be an object-specific scaling, the fluctuation of an event weight within the weight's uncertainty, or a result of a method. Overall, this analysis remains statistics limited; nevertheless, p_T^{miss} reconstruction and the btagging uncertainties have an important impact.

Systematic effects are classified into *rate* and *shape* systematics. Rate systematics affect the overall yield of a process, and are applied equally to all bins in a distribution at limit-setting. The limit setting details are given in Chapter 9. Shape systematics capture bin-by-bin dynamics. They preserve the overall normalization, but allow for the shape of a distribution to fluctuate within uncertainties. The relative differences between the deviated and the nominal RJR mass estimator distributions are taken as a differential rate uncertainty. For background distributions, all systematics are allowed shape variations, but the shape systematic is removed if it does not affect the expected limit based on studies in simulated data. For signal, if no shape dependence is observed within the statistical uncertainty on the signal, no shape uncertainty is applied.

8.1 $Z + \text{jets}$ Estimation with α -Method

Out of the two steps in the α -method, the normalization and the extrapolation via the α ratio, only the extrapolation step contributes uncertainty to the final $Z + \text{jets}$ estimation. The derived normalization is only applied to the Drell-Yan simulation, so the effect cancels in the ratio of Eq. 7.9. Each of the ten fit parameters in the extrapolation (two from each exponential fit) contribute an uncertainty to the final result. The uncertainties on the fit parameters encode the statistical uncertainty of the fitted distribution.

The red one- σ uncertainty bands in Fig. 7.9 represent the “fit envelope”. The fit envelope encloses all of the possible functions generated when the fit parameters are allowed to take their $\chi^2 = \chi_{\text{min}}^2 + 1$ values. This band is an *over*-estimate of the true fit uncertainty, because the parameters and their respective uncertainties are

correlated. Using this uncertainty band also omits the shape variations that characterize the effects of the parameter uncertainties on the fit result. To capture this affect, and to avoid over estimating the uncertainty, a principle component analysis is performed. The fit parameters are decorrelated, individually shifted by their decorrelated uncertainties, and the extrapolation is repeated.

The parameter decorrelation merely rotates the parameter-vector to a basis where the parameters are orthogonal. First we diagonalize the nominal covariance matrix Σ and find the eigenvalues and eigenvectors to form the new decorrelated basis. The covariance matrix of the fit parameters can then be written as

$$\Sigma = \Phi \Lambda \Phi^T,$$

where Φ is a matrix formed from the eigenvectors of Σ and Λ is the diagonal matrix made with the eigenvalues of Σ . In this formulation, Λ is the decorrelated covariance matrix. The original fit parameters X can be transformed into the decorrelated space via

$$Y = \Phi^T X,$$

where Y are the parameters in the decorrelated space. Each parameter Y_i is then shifted by its uncertainty,

$$Y'_i = Y_i \pm \lambda_i^{1/2},$$

where λ is a vector of the eigenvalues of Λ , and Y' is a vector of the decorrelated parameters with only the i th parameter shifted. To propagate the effect through the α -method, we return to the original basis, giving

$$X' = \Phi Y',$$

where the X' is a vector of the parameters representing the shift of a single parameter in the decorrelated space. For each parameter, the uncertainty is derived by propagating a new function that is drawn with the parameters in X' through the rest of the α -method. This produces a total of ten systematic uncertainties.

An additional uncertainty for the fit to sideband data comes from the choice of fit function. The sideband data fit is severely limited by the lack of data in the high RJR Z' mass estimator region. This is not appropriately addressed by extending the simple exponential function to higher values. To have a more conservative estimate of the “down” fluctuations, a more steeply falling alternative fit function of

$$f_{\text{altdwn}}(x) = \exp^{A+Bx+C\sqrt{x}}, \quad (8.1)$$

is applied. Similarly, to have a maximally conservative “upper” alternative function, a function that does not asymptotically approach zero, but instead approaches a constant value, is also applied:

$$f_{\text{altup}}(x) = \exp^{A+Bx} + C. \quad (8.2)$$

The alternative functions are taken as limiting cases with no prior constraints on the distribution in the extrapolation region. We therefore estimate a 68% confidence level assuming a uniform distribution. The shifted function to capture the “up” alternative fit is given as

$$f_{\text{sideband,up}}(x) = 2 \frac{f_{\text{altup}}(x) - f_{\text{nominal}}(x)}{\sqrt{12}} + f_{\text{nominal}}(x), \quad (8.3)$$

and the “down” is similarly parameterized as

$$f_{\text{sideband,down}}(x) = f_{\text{nominal}}(x) - 2 \frac{f_{\text{nominal}}(x) - f_{\text{altdwn}}(x)}{\sqrt{12}}, \quad (8.4)$$

Figures 8.1 and 8.2 show the per-parameter shifted fits (up and down, respectively) that serve as the inputs to the α ratio. Figures 8.3 and 8.4 show the per-parameter up and down shifted fits, respectively, that go into the $t\bar{t}$ and diboson subtracted data sideband function. Finally, Figs. 8.5 and 8.6 show the up and down shifts when propagated through the entire extrapolation, with Figs. 8.7 and 8.8 showing the $t\bar{t}$ and diboson corrected data sideband function and the extrapolation across a larger range in RJR Z' mass estimator. Table 8.1 summarizes the uncertainties for the α -method.

8.2 Data-driven $t\bar{t}$ estimation

The data-driven $t\bar{t}$ estimation method has three intrinsic uncertainties. The first is the statistical uncertainty that comes from the $e\mu$ data which is not treated as a systematic. The next uncertainty is the uncertainty on the $\mu\mu$ to $e\mu$ ratio, which comes from the statistical uncertainty of the Monte Carlo. The third is the shape uncertainty from the difference between the electron and muon reconstruction in CMS. Statistical uncertainties are dominant, so this effect is neglected. The $\mu\mu$ to $e\mu$ ratio uncertainty produces a 3% uncertainty on the final $t\bar{t}$ yield, and is applied as a




Figure 8.1: Left: Drell-Yan signal region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted **up** in the decorrelated space. Right: Same distributions but for the Drell-Yan sideband. These functions are used to build the α ratio. The uncertainties are small, but are noticeable when propagated through the method.




Figure 8.2: Left: Drell-Yan signal region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted **down** in the decorrelated space. Right: Same distributions but for the Drell-Yan sideband. These functions are used to build the α ratio. The uncertainties are small, but are noticeable when propagated through the method.




Figure 8.3: Left: data sideband region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted **up** in the decorrelated space. Middle: Same description but for the $t\bar{t}$ MC contribution. Right: Same description but for the diboson contribution. The $t\bar{t}$ and diboson functions are subtracted from the data sideband distribution to form the distribution that gets multiplied by the α ratio.




Figure 8.4: Left: data sideband region fits for full Run 2. Each additional fit represents the fit result with the labeled parameter shifted **down** in the decorrelated space. Middle: Same description but for the $t\bar{t}$ MC contribution. Right: Same description but for the diboson contribution. The $t\bar{t}$ and diboson functions are subtracted from the data sideband distribution to form the distribution that gets multiplied by the α ratio.




Figure 8.5: Left: Different α ratio that are generated by varying one of the input function's parameters **up**. Middle: Different subtracted data sideband functions generated by varying one of the input function's parameters **up**. Right: The extrapolation to the signal region when one of the input parameters has been varied **up**. See Figs. 8.7 and 8.8 for larger range, logarithmic scale versions of the middle and right plots.




Figure 8.6: Left: Different α ratio that are generated by varying one of the input function's parameters **down**. Middle: Different subtracted data sideband functions generated by varying one of the input function's parameters **down**. Right: The extrapolation to the signal region when one of the input parameters has been varied **down**. See Figs. 8.7 and 8.8 for larger range, logarithmic scale versions of the middle and right plots.




Figure 8.7: Subtracted sideband distributions. Same as the middle plots in Figs. 8.5 and 8.6, but with extended range and logarithmic scale.




Figure 8.8: Shifted extrapolations with log scale and extended range. Same as the right-most plots in Figs. 8.5 and 8.6, but with extended range and logarithmic scale.

Breakdown of α -method uncertainties

Parameter	Application Style	Rate uncertainty (down/up)
DY SB fit par0	rate and shape	1%
DY SB fit par1	rate and shape	1%
DY SR fit par0	rate and shape	2%
DY SR fit par1	rate and shape	3%
Data SB fit par0	rate and shape	1%
Data SB fit par1	rate and shape	9%
Data SB Alternative Fits	rate and shape	0.3/3%
TT SB fit par0	rate and shape	0.3/0.4%
TT SB fit par1	rate and shape	2%
VV SB fit par0	rate and shape	0%
VV SB fit par1	rate and shape	0.2%

Table 8.1: Summary of the affects of data and MC statistics on the α ratio extrapolation.

rate systematic. The final $t\bar{t}$ yield can be affected by MC scale factor uncertainties, but all fall within the statistical uncertainty on the scale, and can be neglected.

8.3 Jet Energy Scale

To account for the uncertainty arising from the Jet Energy Scale (JES) corrections each fat jet four-vector is scaled up and down within the JES uncertainty and the Higgs candidate selections are repeated with the shifted jets. This systematic is fully correlated with the type-1 corrections on $\vec{p}_T^{\text{miss}}$, so the $\vec{p}_T^{\text{miss}}$ is shifted at the same time. Figure 8.9 shows the up, down, and nominal distributions for the background estimations. Three example signal distributions are shown in Fig. 8.10. The JES corrections primarily affect MC, so the $t\bar{t}$ estimation remains unaffected. Changes in the JES shift the acceptance of events with jets or p_T^{miss} at the cut threshold, changing the final yields by rejecting or accepting these borderline events. The similarity in magnitude between the diboson and $Z + \text{jets}$ effects, in spite of the partially data-driven α -method, is because the $t\bar{t}$ and VV distributions that are subtracted from the data sideband fit in the $Z + \text{jets}$ estimation are affected by the JES. The application of the JES uncertainties is summarized in Table 8.2.

Summary of JEC Uncertainty Application		
Process	Application Style	Rate uncertainty (down/up)
Signal	rate	0.1%
Drell Yan	shape and rate	1/5%
$t\bar{t}$	rate	0.3/1%
Diboson	rate	3/2%

Table 8.2: Summary of the effects of Jet Energy Scale corrections up and down shifts. For background processes that only have a “rate” systematic, the inclusion of the shape was tested in the limit framework, and resulted in no change to the limit. For signal, it is clear from Fig. 8.10 that a shape treatment is not appropriate.




Figure 8.9: Plots showing the application of the up/down systematic uncertainties for the JES. The left plot is Z + jets with the full α -method applied, and the plot is the diboson MC.




Figure 8.10: Top: JES up/down shifts for the signal mass point $m_{Z'}$ 2000 m_{N_S} 400 m_{N_S} 200. Middle: JES up/down shifts for signal mass point $m_{Z'}$ 4000 m_{N_D} 800 m_{N_S} 200. Bottom: JES up/down shifts for signal mass point $m_{Z'}$ 5500 m_{N_D} 1800 m_{N_S} 200. All plots are shown in the signal region.

Summary of btag Uncertainty Application		
Process	Application Style	Rate uncertainty (down/up)
Signal	rate	11%
Drell Yan	shape and rate	4%
$t\bar{t}$	rate	0.1%
Diboson	rate and shape	10.5%

Table 8.3: Summary of the effects of btagging up and down shifts. For signal, it is clear from Fig. 8.12 that a shape treatment is not appropriate, and the dominant behavior is captured in the histogram normalization.

8.4 b-tagging

Systematic uncertainties due to b-tagging are derived by shifting the b-tagging event weight up and down by the scale factor's uncertainty. Figure 8.11 shows the up, down, and nominal distributions for the respective backgrounds. Figure 8.12 shows the up/down shifts for three example signal points. Table 8.3 summarizes the derived uncertainties for the DeepAK8MD tagging. Unlike the JES corrections in Section 8.3, the partially data-driven α -method does reduce the impact of the b-tagging systematic uncertainty for the Z + jets background. Changing the event weight does not change the event's acceptance, in contrast to the JES uncertainty. When events cannot shift analysis region, the normalization to the sideband data can absorb the btag scale factors' effects. The differences in response to the b-tagging uncertainties between the data-driven and pure MC samples is an indicator of the benefits of data-driven methods.




Figure 8.11: Plots showing the application of the up/down systematic uncertainties for b-tagging. The left plot is Z + jets with the full α -method applied, and the plot is the diboson MC.




Figure 8.12: Top: btag up/down shifts for the signal mass point $m_{Z'}$ 2000 m_{N_S} 400 m_{N_S} 200. Middle: btag up/down shifts for signal mass point $m_{Z'}$ 4000 m_{N_D} 800 m_{N_S} 200. Bottom: btag up/down shifts for signal mass point $m_{Z'}$ 5500 m_{N_D} 1800 m_{N_S} 200. All plots are shown in the signal region.

8.5 p_T^{miss} Uncertainties

Since $\vec{p}_T^{\text{miss}}$ is a quantity derived from the kinematics of the visible particles in an event, any uncertainty with visible particles must be propagated to the $\vec{p}_T^{\text{miss}}$. The type-1 corrections are fully correlated with the JES uncertainties on the jets and are covered in Section 8.3. The main uncertainty on $\vec{p}_T^{\text{miss}}$ in this analysis is the “unclustered MET.” Unclustered MET (missing transverse energy) is the pseudo-object built by vectorially summing the PF candidates in an event that are not clustered into a jet. The MET and p_T^{miss} are the same: the term “MET” simply implies a massless invisible particle. To generate the up and down distributions, the nominal $\vec{p}_T^{\text{miss}}$ is fluctuated by the uncertainty on the unclustered components, and the selections are redone. Figure 8.13 shows the up, down, and nominal background distributions. Figure 8.14 shows the up/down distributions for three example signal points. The application of the unclustered MET uncertainties is summarized in Table 8.4. Unclustered MET systematics mirror the relative effects of the JES systematics because the unclustered MET fluctuations also change the number of accepted events.

Summary of unclustered MET Uncertainty Application		
Process	Application Style	Rate uncertainty (down/up)
Signal	rate	0.1%
Drell Yan	shape and rate	5/10%
$t\bar{t}$	rate	0.5/0.8%
Diboson	rate	9/8%

Table 8.4: Summary of the effects of unclustered MET up and down shifts. For signal, it is clear from Fig. 8.14 that a shape treatment is not appropriate, and the dominant behavior is captured in the histogram normalization.

8.6 Muon ID

The uncertainties associated with muon identification are captured with the uncertainties on the muon ID scale factors. For the dimuon reconstruction of the Z candidate, the associated muon ID event weight is the product of the scale factors of each of the constituent muons. The uncertainty on this weight is then given by,

$$\sigma_{\text{muonID}} = sf_{\mu 1} sf_{\mu 2} \sqrt{\frac{\sigma_{sf_{\mu 1}}^2}{sf_{\mu 1}} + \frac{\sigma_{sf_{\mu 2}}^2}{sf_{\mu 2}}}.$$

The up and down distributions are generated by varying the muon ID scale factor




Figure 8.13: Plots showing the application of the up/down systematic uncertainties for unclustered MET. The left plot is $Z + \text{jets}$ with the full α -method applied, and the plot is the diboson MC.

event weight up and down by the derived uncertainty. Overall, this ends up being a variation of order 1%. Only rate systematics are considered.

8.7 Muon Trigger Scale Factor

Muon trigger scale factors capture the uncertainty on the muon trigger efficiency in Monte Carlo. The trigger event weight given by Eq. 6.4 is dependent on four parameters, each of which have an associated uncertainty. Ignoring the correlations between the scale factor uncertainties and the efficiency measurement uncertainties, the uncertainty on the muon trigger scale factor event weight is

$$\sigma_{sf_{event}} = \sqrt{\left(\frac{\partial sf_{event}}{\partial sf_{\mu 1}}\right)^2 \sigma_{sf_{\mu 1}}^2 + \left(\frac{\partial sf_{event}}{\partial sf_{\mu 2}}\right)^2 \sigma_{sf_{\mu 2}}^2 + \left(\frac{\partial sf_{event}}{\partial eff_{\mu 1}}\right)^2 \sigma_{eff_{\mu 1}}^2 + \left(\frac{\partial sf_{event}}{\partial eff_{\mu 2}}\right)^2 \sigma_{eff_{\mu 2}}^2}. \quad (8.5)$$

The efficiency uncertainties are statistical only and are very small in most bins. These are further suppressed when squared, allowing Eq. 8.5 to be reduced to

$$\sigma_{sf_{event}} = \sqrt{\left(\frac{\partial sf_{event}}{\partial sf_{\mu 1}}\right)^2 \sigma_{sf_{\mu 1}}^2 + \left(\frac{\partial sf_{event}}{\partial sf_{\mu 2}}\right)^2 \sigma_{sf_{\mu 2}}^2}. \quad (8.6)$$

Shifting the muon trigger event weights up and down by equation 8.6 produces about a 0.1% effect on the final yields.




Figure 8.14: Top: Unclustered MET up/down shifts for the signal mass point $m_{Z'}$ 2000 m_{N_S} 400 m_{N_S} 200. Middle: Unclustered MET up/down shifts for signal mass point $m_{Z'}$ 4000 m_{N_D} 800 m_{N_S} 200. Bottom: Unclustered MET up/down shifts for signal mass point $m_{Z'}$ 5500 m_{N_D} 1800 m_{N_S} 200. All plots are shown in the signal region.

8.8 HEM15/16

The HEM15/16 incident is the loss of 40° in ϕ of the HCAL Minus Endcap (HEM) in the second half of 2018 data-taking. HEM15 and HEM16 were lost due to a malfunction in the front end CAEN power supplies after an emergency shutdown triggered by a false fire alarm in the CMS cavern. This gap in the HCAL endcap affects forward jet energy measurements, p_T^{miss} measurements, and electron identification. To characterize the impact, one scales-down the jet four-vectors in the affected regions in MC between 20 and 30% depending on the η position. The scaling is also propagated to the $\vec{p}_T^{\text{miss}}$.

Since this analysis uses only barrel-jets and has a moderately high p_T^{miss} requirement, we do not expect sensitivity to the HEM15/16 loss. For a simple check, we scale the entire 2018 integrated luminosity in MC. The true effect on 2018 is 65% of this study's magnitude. Figure 8.15 shows the MC distributions used for the α ratio background estimation for the total 2018 luminosity. For the overall method, 2018, 2017, and 2016 luminosities are combined, reducing the observed percent-level effects to sub-percent level effects. This falls within the uncertainty on the individual fits. Therefore this systematic is neglected for the $Z + \text{jets}$ estimation. For diboson processes, we use a fully MC derived background. Figure 8.16 shows the 2018 signal region shapes when the entire 2018 luminosity is scaled. The two diboson processes are treated together, and represent roughly 2% of the entire signal region when all three years are combined. A 2% HEM15/16 effect, that occurs in less than one-third of the MC, that in total makes up 2% of the signal region, is neglected.

Since the HEM15/16 loss also affects the reconstruction of electrons, the data-driven $t\bar{t}$ background taken from the $e\mu$ data set requires special consideration. The effect of the loss on electron reconstruction is solved by vetoing events with an electron in that region, and by weighting these events in MC by the fraction of luminosity where they are valid (36% of 2018 data). The jet scaling effect exists in addition to the electron effect. In the data-driven $t\bar{t}$ background, only the comparative yields in MC between the $\mu\mu$ and $e\mu$ data sets matter to derive the normalization difference between the channels. Table 8.5 summarizes the various HEM effects on these yields. The effects fall well within the statistical uncertainty of the Monte Carlo samples used to derive the ratio, even with the jet effects over-estimated as with the diboson and α -method tests. We neglect the HEM effects in the uncertainty on the $\mu\mu$ to $e\mu$ ratio.

The jet energy scaling could also affect the signal. The scaling effects on three signal mass points are shown in Fig. 8.17. We observe negligible effects.




Figure 8.15: 2018 MC input distributions to the alpha method for the $Z + \text{jets}$ estimation. Entire 2018 luminosity scaled, so any deviations over-estimate their expected value. Deviations are well within the fit uncertainties for the method, especially once combined with 2016 and 2017, as is done.




Figure 8.16: 2018 MC input distributions for the diboson background. Entire 2018 luminosity scaled, so any deviations over-estimate their expected value. Deviations are negligible, especially once combined with 2016 and 2017, as is done.

yield description	yield and statistical uncertainty in total mass region at Preselection
Nominal $\mu\mu$ MC	146.50 ± 4.20
HEM jet scaling $\mu\mu$ MC	145.64 ± 4.19
$e\mu$ no veto, no scale MC	140.11 ± 5.89
$e\mu$ veto MC	139.62 ± 5.88
$e\mu$ veto, HEM jet scale MC	138.54 ± 5.85


Table 8.5: $t\bar{t}$ estimation HEM effects on input yields.

Figure 8.17: Observable signal distributions in the dimuon channel with the HEM scaling applied. Entire sample scaled, so any deviations over-estimate their expected value. Deviations are negligible, especially once combined with 2016 and 2017, as is done.

8.9 Parton Distribution Function (PDF) Scale

There is an uncertainty associated with the scale used for the Parton Distribution Function (PDF) definition for Monte Carlo. For each MC event, there are at least 100 PDF replicas stored, each replica representing an event weight following the Gaussian distribution of the PDF scale. A reasonable uncertainty to take is the standard deviation of the replica distribution. The 100 replica weights are sorted in ascending order, and the uncertainty is taken as

$$\sigma_{\text{pdf}} = \frac{\text{Replica}_{84} - \text{Replica}_{16}}{2}. \quad (8.7)$$

Assuming a nominal weight of 1, the up and down distributions are generated by scaling the MC events with an event weight of $1 \pm \sigma_{\text{pdf}}$. Table 8.6 summarizes the effects of the PDF scales.

Summary of PDF Scale Uncertainty Application		
Process	Application Style	Rate uncertainty (down/up)
Signal	rate	2-25%
Drell Yan	rate	1%
$t\bar{t}$	rate	0.1%
Diboson	rate	2%

Table 8.6: Summary of the effects on signal region yield of pdf scale shifts. Signal effects get larger with Z' mass.

8.10 QCD renormalization and Factorization scale

For each Monte Carlo event, there are event weights that represent the effects of shifting the QCD renormalization and factorization scales used to generate the MC. To get a conservative estimate of the effect, the largest and smallest event weights are taken and applied to form the up and down shifted distributions. Table 8.7 summarizes the differences with respect to the nominal yields with this prescription.

Summary of QCD Scale Uncertainty Application		
Process	Application Style	Rate uncertainty (down/up)
Signal	rate	10-20%
Drell Yan	rate	11%
$t\bar{t}$	rate	0.3%
Diboson	rate	30%

Table 8.7: Summary of the effects on signal region yield of maximal qcd scale shifts. Signal effects get larger with Z' mass.

Chapter 9

Statistical Interpretation and Results

A search for BSM physics quantifies how *unlike* the observed data is from the expectations of the standard model. To perform a search, a test statistic is used that determines how unlikely it is that the observed data is due to SM processes. Without a clear discovery, a “limit” can be set on the production rate of new physics processes. A limit captures the sensitivity of a search: the lowest rate of production that could be reasonably measured. To do this, we prepare two hypotheses. The first is the background-only (SM) hypothesis, built from the background model detailed in Chapter 7. The second is the signal-plus-background hypothesis, which combines the derived background expectation and the leptophobic Z' signal model. To compare them, a binned profile likelihood method is used.

This chapter presents the expected frequentist limits in the asymptotic approximation for the production cross section of the leptophobic Z' decaying to neutral anomalous. These limits characterize the sensitivity of this analysis. Section 9.1 describes the statistical analysis and Section 9.2 details the input model design and structure. Since limiting-setting relies on numerical fitting, Section 9.3 summarizes the statistical tests that verify the analytical power and robustness of the model. Finally, Section 9.4 presents the expected limits which characterize the sensitivity of this analysis.

9.1 Statistical Analysis

The statistical analysis uses a binned profile likelihood ratio [Cowan_2011]. The likelihood function $L(\mu, \theta)$ is the product of the Poisson probabilities to observe a signal strength μ given the set of nuisance parameters θ in each bin of the analysis. The signal strength μ is a multiplier of the production cross section used in the input model and the nuisance parameters θ capture the effects of the systematic uncertainties. The profile likelihood ratio is defined as

$$\lambda(\mu) = \frac{L(\mu, \hat{\hat{\theta}})}{L(\hat{\mu}, \hat{\hat{\theta}})}, \quad (9.1)$$

where $\hat{\theta}$ are the values of θ that maximize L for the specified μ . The $\hat{\mu}$ and $\hat{\theta}$ are the values that generically maximize L for the observed data. A $\lambda(\mu)$ close to unity implies good agreement between a predicted μ and observed $\hat{\mu}$. The value of $\hat{\mu}$ is determined from a maximum likelihood fit to the data comprising both signal and background models. We can maximize Eq. 9.1 to extract the signal strength that best describes the data, or equivalently, we can minimize. In this case, a new test statistic,

$$q_\mu = -2\ln(\lambda(\mu)) , \quad (9.2)$$

where q_μ is negative logarithm of the profile likelihood ratio, is used.

The significance of an observed excess of events over the background expectation is calculated using the test statistic q_0 defined as

$$q_0 = \begin{cases} -2\ln\lambda(0) & \hat{\mu} \geq 0 \\ 0 & \hat{\mu} < 0 \end{cases} , \quad (9.3)$$

where $\lambda(0)$ represents the likelihood ratio where $\mu = 0$. The $\mu = 0$ case is the null hypothesis: the absence of detectable signal. To set an exclusion limit, we modify the q_μ test statistic,

$$q_\mu = \begin{cases} -2\ln \frac{L(\mu, \hat{\theta})}{L(0, \hat{\theta}(0))} & \hat{\mu} < 0 \\ -2\ln \frac{L(\mu, \hat{\theta})}{L(\hat{\mu}, \hat{\theta})} & 0 \leq \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} . \quad (9.4)$$

The test statistic goes to zero when $\hat{\mu} > \mu$. In the context of *excluding* values of μ from possibility, having a measured $\hat{\mu} > \mu$ does not necessarily imply incompatibility with data, but instead implies incompatibility with μ being the lower limit.

The p -value p_μ quantifies the level of agreement between the measurement and the hypothesized μ ,

$$p_\mu = \int_{q_{\hat{\mu}}}^{\infty} f(q_\mu|\mu) dq_\mu , \quad (9.5)$$

where $f(q_\mu|\mu)$ is the probability density function (PDF) of q_μ assuming μ . The p -values are used to calculate the upper limit using the “ CL_s criterion.” For the CL_s upper limit, a signal hypothesis is considered excluded at the 95% confidence level if

$$\frac{p_\mu}{1 - p_{\text{bkg}}} \leq 0.05, \quad (9.6)$$

where p_{bkg} is the p -value under the background-only hypothesis. With large yields, $\hat{\mu}$ follows a Gaussian with mean μ' and standard deviation σ . Within this limit, one can write

$$-2\ln(\lambda(\mu)) = \frac{(\mu - \hat{\mu})^2}{\sigma^2} + \mathcal{O}(1/\sqrt{N}), \quad (9.7)$$

where N is the data sample size. Assuming $\mu' = \mu$, the p -value calculation becomes

$$p_\mu = \begin{cases} 1 - \Phi(\sqrt{q_\mu}) & 0 < q_\mu < \mu^2/\sigma^2 \\ 1 - \Phi(\frac{q_\mu + \mu^2/\sigma^2}{2\mu/\sigma}) & q_\mu > \mu^2/\sigma^2 \end{cases}, \quad (9.8)$$

where Φ is the cumulative distribution function of a standard Gaussian. This is called the asymptotic approximation. Outside of the asymptotic approximation, Monte Carlo techniques are needed to build $f(q_\mu|\mu)$ and to calculate σ . In the asymptotic limit, the corresponding parameters can be pulled directly from the Asimov data set. The Asimov data set is a representative data set built from the expectation values of the parameters in the likelihood. It is defined as the data set where the estimations of the parameter values are the true parameter values [Cowan_2011]. This set is built by finding the data set where the partial derivatives of the likelihood function with respect to the nuisance parameters are equal to zero.

9.2 Input Model

The input model consists of 8 bins of the RJR Z' mass estimator spectrum. The first 7 bins cover the mass estimator range between 1400 and 2800 GeV in 200 GeV divisions. The final bin is an “overflow” bin that extends the range from 2800 to 5000. This 1400 to 5000 range is chosen to avoid the turn-on curve of the backgrounds due to trigger and reconstruction, and to balance the deterioration of the background estimations with maintaining signal acceptance in the high estimator region.

The validity of the background estimate ultimately depends on the statistics available in the data and MC sidebands. The relationship between the RJR mass estimator’s input kinematics and the sideband tails governs the input model range. Figure 9.1 shows the relationship between the mass estimator for the highest $m_{Z'}$ mass point and a few of its kinematic inputs. The highest p_T fat jet in the preselection sideband is around 900 GeV, the highest Z candidate p_T is past 800 GeV, and the highest MET




Figure 9.1: Relationship between the RJR Z' mass estimator and select kinematic inputs. Relationship is shown for the $m_{Z'}$ 5500, m_{N_D} 1800, and m_{N_S} 200 signal mass point in the signal region. No cross section scaling is applied.

event is around 400 GeV. The p_T ranges of the scale factors are more restrictive, with 700 GeV for the DeepAK8 score, 500 GeV for the muon trigger scale factors, and 120 GeV for the muon ID scale factors. Drawing boundaries at these points on the corresponding plots in Fig. 9.1 disfavors the generation of RJR mass estimator values above 5000 GeV.

Figures 9.2 and 9.3 show the final binning and statistical uncertainties for the $t\bar{t}$ and diboson backgrounds, respectively. Poisson statistics are used for the $t\bar{t}$, producing one-sided uncertainties for the bins without counts. All function extrapolations described in Section 8.1 are cast into histograms prior to the re-binning, taking the value of the function at the bin center as the expected count. This preserves the validity of the re-binned shapes and yields for the estimations. Finally, Fig. 9.4 shows the signal distribution with the highest masses.




Figure 9.2: The $t\bar{t}$ background estimation, taken from the electron and muon final-state $t\bar{t}$ data.




Figure 9.3: The diboson background estimation, taken from Monte Carlo.




Figure 9.4: Highest $m_{Z'}$ and m_{ND} signal mass point with the final statistical model binning.

9.2.1 Signal Interpolation

The generated signal grid in Fig. 6.3 provides the signal models for the signal plus background likelihood fits. Both the Z' and the N_D masses affect the final state kinematics, so the grid for Monte Carlo generation was chosen specifically to allow for interpolation in the m_{N_D} and $m_{Z'}$ plane with each interpolated point having four adjacent points of full simulation. A Radial Basis Function interpolation [91] is used to estimate the signal shapes at the mass points that lack generated simulation.

A Radial Basis Function (RBF) is any function whose value is only dependent on its distance from a given point. A collection of RBFs relating the points in a data set can define a space (a *basis*) and can therefore be used to interpolate values between the existing points. The missing signal RJR Z' mass estimator shapes in our analysis are built from values interpolated on planes in m_{N_D} and $m_{Z'}$ space at specific quantiles of the simulated signal shapes. Figure 9.5 shows a cartoon of the quantiles, and the resulting plane. To build the shapes, planes of the generated signal shape values at 10,000 equally spaced quantiles are built, each in turn providing a basis to interpolate the intermediate missing values in the m_{N_D} and $m_{Z'}$ plane. In our interpolation a thin plate spline function is used as the RBF. The overall normalization is also derived using RBF interpolation, with the set of generated MC normalizations providing the input data set.

Figure 9.6 shows the Monte Carlo generated and interpolated RJR Z' mass estimator shapes for the signal at both fixed m_{N_D} and fixed Z' , for $m_{N_S} = 200$ GeV. The interpolated shapes behave well with respect to the generated Monte Carlo.

9.3 Checks of Statistical Model

To exercise the limit framework, of three benchmark mass points are used,

1. $m_{Z'} 2000, m_{N_D} 400, m_{N_S} 200,$
2. $m_{Z'} 4000, m_{N_D} 800, m_{N_S} 200,$
3. $m_{Z'} 5500, m_{N_D} 1800, m_{N_S} 200,$

because together they span the generated signal grid given in Section 6.2. Each of these represent different kinematic regimes. The first mass point has the RJR Z' mass estimator peak in the high background region of the search window, and the m_{N_D} and m_{N_S} are boosted, since m_{N_D} has a value far away from half of the Z' value. The second and final mass points both have Z' mass estimators peaking in the tail of the distribution, with various levels of boost in the final state. Both N_D particles




Figure 9.5: Cartoon of the relationship between the q_{ith} quantiles and the plane in $mass_1 mass_2$ -plane used to interpolate the quantile position for new mass points. In the figure, $mass_1$ and $mass_2$ are arbitrary stand-ins for $m_{Z'}$ and m_{N_S} in our signal model.

are produced near the $1/2m_{Z'}$ threshold, putting them nearly at rest. Figure 9.7 illustrates these three signal region scenarios.

9.3.1 Pulls

The systematics described in Chapter 8, as well as the statistical uncertainties that come from the backgrounds and signal, all enter in the limit calculation as nuisance parameters. The systematic studies proceeding this section inform the “pre-fit” values of these parameters. Once the likelihood fit to data is preformed to set the limit (or quantify the significance of an excess), the nuisance parameters take on new values to minimize Eq. 9.4.

To understand this progression, nuisance parameter “pulls” are generated with 500 toy data pseudoexperiments with expected signal strengths of either 0 or the median expected limit. The “pull” is defined as

$$pull = \frac{\theta - \theta_I}{\sigma_I},$$

where θ is the post-fit nuisance parameter, θ_I is the pre-fit nuisance parameter, and σ_I is the pre-fit uncertainty on the nuisance parameter. Figures 9.8 through 9.10




Figure 9.6: Selected values of fixed m_{N_D} and $m_{Z'}$ comparing the Monte Carlo generated signal samples and the interpolated samples (bold lines). The left column is at fixed m_{N_D} , and the right column is at fixed $m_{Z'}$. All plots are for the $m_{N_S} = 200$ GeV case. The interpolation appears to be correctly modeling the shape and normalization trends.




Figure 9.7: Signal region distributions for the example signal points used to test the statistical model.

show the pulls with an injected signal equal to the median expected limit. Figure 9.11 shows the background only pull for the $m_{Z'}$ 5500 m_{N_D} 1800 m_{N_S} 200 GeV mass point. The $t\bar{t}$ statistical nuisance parameters have the largest pulls, because they are most likely to be affected by statistical fluctuations in the pseudoexperiments.

9.3.2 Impacts

Impact plots show the importance of a specific nuisance parameter in the likelihood fit. Impacts are quantified by measuring the change in the derived signal strength when a specific nuisance parameter is shifted to its $\pm 1\sigma$ post-fit value, the others are frozen at their nominal values, and the fit is repeated.

Figures 9.12 and 9.13 show the impact plots of the background only fits for the $m_{Z'}$ 2000 m_{N_D} 400 m_{N_S} 200 GeV and $m_{Z'}$ 5500 m_{N_D} 1800 m_{N_S} 200 GeV signal mass points, respectively. As expected, the low statistics of the $t\bar{t}$ estimation in the expected signal range dominate. The differences in the signal peak location are evident in the impacts: the lower-bin statistics dominate in the light signal case, and the higher bins dominate in the heavy sample. The Z + jets contribution in the high-tails is also reflected in the importance of the α -method systematics in the tails.

9.3.3 Signal Injection Tests

The signal injection tests verify the sensitivity to the signal model and test that the model is not biased to give the incorrect signal strength. Table 9.1 summarizes the four injected signal strengths that correspond with 0 expected signal, the median expected limit, and finally the lower and upper expected limits.

For each mass point and injected signal strength, 500 pseudoexperiments are generated. Figure 9.14 shows the linearity of the measured signal strength versus the




Figure 9.8: Nuisance parameter pulls for fits to 500 toys with an expected signal strength equal to the median expected limit for mass point $m_{Z'}$ 2000 m_{N_D} 400 m_{N_S} 200 GeV. Pre-Fit line indicates the bare nuisance parameter value, not a pull.




Figure 9.9: Nuisance parameter pulls for fits to 500 toys with an expected signal strength of 0 for mass point $m_{Z'}$ 4000 m_{N_D} 800 m_{N_S} 200 GeV. Pre-Fit line indicates the bare nuisance parameter value, not a pull.




Figure 9.10: Nuisance parameter pulls for fits to 500 toys with an expected signal strength of 0 for mass point $m_{Z'} \ 5500 \ m_{N_D} \ 1800 \ m_{N_S} \ 200 \text{ GeV}$. Pre-Fit line indicates the bare nuisance parameter value, not a pull.




Figure 9.11: Nuisance parameter pulls for fits to 500 toys with an expected signal strength equal to the median expected limit for mass point $m_{Z'} = 5500$ m_{N_D} 1800 m_{N_S} 200 GeV. Pre-Fit line indicates the bare nuisance parameter value, not a pull.




Figure 9.12: Impact of nuisance parameters for the background only hypothesis with a signal hypothesis of $m_{Z'} 2000 \text{ GeV}$, $m_{N_D} 400 \text{ GeV}$, $m_{N_S} 200 \text{ GeV}$. The Δr is the change in the extracted signal strength with respect to the nominal when the nuisance parameter is set to its $\pm 1\sigma$ value and the fit repeated.




Figure 9.13: Impact of nuisance parameters for the background only hypothesis with a signal hypothesis of $m_{Z'}$ 5500 GeV, m_{N_D} 1800 GeV, m_{N_S} 200 GeV. The Δr is the change in the extracted signal strength with respect to the nominal when the nuisance parameter is set to its $\pm 1\sigma$ value and the fit repeated.

Mass Point	Injected Signal Strengths (fb)
$m_{Z'} = 2000 \, m_{N_D} = 400 \, m_{N_S} = 200$	0, 0.53, 0.79, 1.20
$m_{Z'} = 4000 \, m_{N_D} = 800 \, m_{N_S} = 200$	0, 0.19, 0.28, 0.45
$m_{Z'} = 5500 \, m_{N_D} = 1800 \, m_{N_S} = 200$	0, 0.24, 0.38, 0.65

Table 9.1: Signal Injection Test Parameters.




Figure 9.14: Trend of signal strength (r) measured versus the signal strength injected. Left: signal sample with $m_{Z'} = 2000, m_{N_D} = 400, m_{N_S} = 200$. Middle: signal sample with $m_{Z'} = 4000, m_{N_D} = 800, m_{N_S} = 200$. Right signal sample $m_{Z'} = 5500, m_{N_D} = 1800, m_{N_S} = 200$.

injected signal strength for each mass point. A constant bias is observed, but this is due to the low background expectation. Our signal predominantly exists in the tails of the background distribution where the background expected is exceedingly low. In this test the signal strength is allowed to take nonphysical values and the signal+background can converge with a negative extracted signal strength. This results in a slight “skew”, or excess of events in the lower side of the signal injection test distributions, biasing the result. This was observed in previous limit variations and Appendix D contains investigations showing that this behavior is only due to the limited statistics. A linear correction can be applied in this case.

Figures 9.15 through 9.17 show the difference between the measured signal strength and the injected signal strength and the pull of the signal strength. The measured and injected signal strength difference should be centered at 0. The pull is defined as the difference between the measured and the injected signal strength, divided by the uncertainty on the measured strength defined by

$$\sigma_{hilow} = \sigma_{r,high}((r - r_{inj}) < 0) + \sigma_{r,low}((r - r_{inj}) > 0),$$

where $\sigma_{r,high}$ and $\sigma_{r,low}$ are the upper and lower uncertainties on the measured signal strength. After plotting, each distribution is fit with a simple Gaussian. The biases are as expected for low-background analyses.




Figure 9.15: Differences (left) and pulls (right) for the measured and expected signal strengths for $m_{Z'} = 2000$, $m_{N_D} = 400$, $m_{N_S} = 200$. First row has an injected signal of 0, the second has the lower expected limit, the third the median expected limit, and the final the upper expected limit.




Figure 9.16: Differences (left) and pulls (right) for the measured and expected signal strengths for $m_{Z'} = 4000$, $m_{N_D} = 800$, $m_{N_S} = 200$. First row has an injected signal of 0, the second has the lower expected limit, the third the median expected limit, and the final the upper expected limit.




Figure 9.17: Differences (left) and pulls (right) for the measured and expected signal strengths for $m_{Z'} = 5500$, $m_{N_D} = 1800$, $m_{N_S} = 200$. First row has an injected signal of 0, the second has the lower expected limit, the third the median expected limit, and the final the upper expected limit.

Progression of limit changes		
Signal Mass Point (GeV) ($m_{Z'}$, m_{N_D} , m_{N_S})	MC $t\bar{t}$	$e\mu$ $t\bar{t}$
(2000, 400, 200)	1.08 fb	0.81 fb
(4000, 800, 200)	1.27 fb	1.25 fb
(5500, 1800, 200)	1.74 fb	1.78 fb

Table 9.2: Limit progression applying the data-driven $t\bar{t}$ background estimation method.

9.4 Results

We report the expected sensitivity of the analysis in terms of the median expected upper limits on the cross section times branching fraction times acceptance of our various leptophobic Z' scenarios. The branching fraction is assumed to be one. The grid of expected fiducial cross section upper limits in the $m_{Z'}$ and m_{N_D} mass plane for the $m_{N_S} = 200$ GeV and $m_{N_S} = 1$ GeV final states are presented in Fig. 9.18 and Fig. 9.19, respectively. The analysis loses sensitivity in the extremes of the Z' mass where either background processes dominate or the signal rate is limited by the available center of mass collision energy. Figures 9.20 and 9.21 show the limit projections along the m_{N_D} axis for a given $m_{Z'}$ for the $m_{N_S} = 200$ GeV and $m_{N_S} = 1$ GeV final states, respectively.

The lower mass ranges for the $m_{N_S} = 200$ GeV are limited by the boost of the final state particles; without a heavy Z' , the H is not boosted enough to satisfy the fat jet requirements. The higher Z' masses also have lower sensitivity than those in the mid-range because larger fractions of the RJR Z' mass estimator peak exist outside of the limit-model acceptance. A slight concave-down bowing is observed in the higher Z' masses. The two m_{N_D} extremes represent the two scenarios introduced in Fig. 6.1. In the lighter N_D scenario, the N_D particle itself is boosted, resulting in a back-to-back boosted final state. These events definitively satisfy the jet boost requirement, but lack $\vec{p}_T^{\text{miss}}$. The opposite scenario has the N_D produced almost at rest, and the decay produces isotropic distributions of final state particles with high p_T^{miss} . Increasing the m_{N_D} transitions between the jet momentum and p_T^{miss} dominated final states creating a region where neither requirement is met.

Overall, the sensitivity of the $m_{N_S} = 1$ GeV final state is comparable to the $m_{N_S} = 200$ GeV case. The analysis is not optimized for the $m_{N_S} = 1$ GeV, and slightly different behavior is observed. Sensitivity is lost for the high mass Z' with a large mass splitting because of the extreme back-to-back topology of the final state.

The data-driven $t\bar{t}$ estimation dramatically increases the sensitivity for the low $m_{Z'}$ mass signals by removing the impact of $t\bar{t}$ MC systematics. See Table 9.2.




Figure 9.18: Asymptotic fiducial cross section upper limits for 2016-2018 are shown in the m_{N_D} and $m_{Z'}$ plane for the $m_{N_S} = 200$ GeV scenario. Limits shown in fb.




Figure 9.19: Asymptotic fiducial cross section upper limits for 2016-2018 are shown in the m_{N_D} and $m_{Z'}$ plane for the $m_{N_S} = 1$ GeV scenario. Interpolated signals are not included. Limits shown in fb.




Figure 9.20: Asymptotic fiducial cross section upper limits for 2016-2018 as a function of m_{N_D} for fixed values of $m_{Z'}$ for the $m_{N_S} = 200$ GeV scenario. Limits shown in fb.




Figure 9.21: Asymptotic fiducial cross section upper limits for 2016-2018 as a function of m_{ND} for fixed values of $m_{Z'}$ for the $m_{N_S} = 1$ GeV scenario. Limits shown in fb.

Chapter 10

Conclusions and Outlook

This thesis has presented a novel search for a leptophobic Z' and selected stories of the CMS HCAL Upgrade. The search presented not only establishes a technique to determine the most sensitive expected exclusion limits for a high-mass leptophobic Z' to date, but pioneers the use of new observables in CMS in an unexplored region of phase space. The work on the Phase 1 Upgrade not only improves the existing performance of CMS detector, but adds new capabilities to prepare CMS for the HL-LHC. We are in an unprecedented era of high energy particle physics at colliders. Without clear theoretical direction for new collider-based physics and with the HL-LHC on the horizon, we need new tools and approaches to exhaust the discovery potential of the machines.

Figure 10.1 shows the most recent ATLAS and CMS dijet resonance limits. These can be interpreted to set exclusion limits on the leptophobic Z' hadronic decay channel. Our expected fiducial cross section limits shown in Fig. 9.18 are more sensitive than both limits across the entire phase space considered in this analysis. Not only does this analysis introduce a new and more sensitive class of models to the LHC search menu, it introduces Recursive Jigsaw Reconstruction to CMS for searches with p_T^{miss} in the final state. Because of this work, RJR has entered the CMS tool-kit to be incorporated into future searches.

Since we are not yet unblinded, we have the freedom to add additional parameters of interest to our search. To align with dijet results, in our next iteration we can extract and set a limit on the Z' coupling. We will also add additional signal interpretations, and plan to include the $Z \rightarrow ee$ channel.

The data-driven background methods developed in this analysis can be used to improve the sensitivity of searches with similar backgrounds. Searches in CMS with high p_T^{miss} requirements and significant jet use are generally heavily impacted by JEC, btag, and p_T^{miss} systematics, but this analysis is an exception to this rule due to our data-driven techniques.

The most natural future extension of the search is to focus on low-mass leptophobic Z' models. With the order-of-magnitude gain over traditional dijet limits in the heavy regime, the anomalous final-state tempts a low mass search. Direct low-mass dijet searches are limited by jet-trigger thresholds. To avoid this, low-mass limits on the




Figure 10.1: Left: Latest CMS Experiment narrow-resonance dijet limits [36]. Right: Latest ATLAS Experiment narrow-resonance dijet limits on excited quark q^* models [92]. Other signal interpretations have similar limits. These results can be interpreted to set limits on the hadronic channel of a leptophobic Z' .

leptophobic Z' coupling exclusively to quarks are found via data-scouting or associated production techniques [93, 94]. The anomalon channel offers a complimentary search strategy without the complications from the jet trigger threshold a priori, and does not suffer from the large QCD backgrounds. Such an analysis would no longer be in the boosted regime, but instead would rely on resolving the two b-jets individually. This future analysis with lower-energy resolved-jets in the final state will benefit from the hardware-level pile-up mitigation now possible with the HCAL upgrade.

With the beginning of Run 3 upon us, the new HCAL front end electronics performance is being explored. So far, all systems perform as expected, though the ramp to maximum luminosity is still underway. At maximum pile-up, we can gauge the full performance of the new segmentation. Unfortunately, RSSI drift has been observed in a few modules in the HCAL Barrel with full-time operation. Two cases have required the switch to the back-up link. While the 2022 year end technical stop will verify the cause, we anticipate that this is due to assembly errors, and that these two suspect VTRxs lack cooling fins. This assembly oversight was rectified for the rework of the HE ngCCMs. There are three other modules that exhibit drift at levels below the instability threshold. While undesirable, this is not unexpected due to the probabilistic nature of the effect. Even with fins, there will still be VTRxs where the effect is not fully mitigated; none-the-less, the CMS HCAL Barrel is expected to operate stably.

While CMS was not originally designed for non-prompt new physics, upgrades like

the HCAL barrel electronics extend CMS's capabilities to entirely new regimes. Like cascade decays, long lived particles (LLPs) are a class of new physics searches that offer unexplored phase space. With timing information now available in the upgraded HCAL read out, Run 3 analyses can use the newly developed LLP triggers. This lays the groundwork for searches that can be further enhanced at the HL-LHC when combined with the upgraded ECAL Barrel timing and the entirely new MIP (minimum ionizing particle) Timing Detector in CMS.

The Phase 1 upgrade to the CMS HCAL ensures that the calorimeter will operate for the foreseeable future. While the endcaps will be replaced before the start of the HL-LHC, the barrel will last the lifetime of CMS. The skills forged in the LHC and HL-LHC upgrades will be useful when building the next generation of detectors.

In searches for new phenomena, we must cast our net wide and deep. Current and future detectors must be designed – or re-designed – to focus on more exotic signals. More complex scenarios can be hiding in LHC data, and it is our job to develop the tools to extract these elusive signatures. The complex final state in this analysis breathes new life into old models and offers new and exciting topologies to explore.

Bibliography

- [1] CMS Collaboration. “Search for resonant and nonresonant new phenomena in high-mass dilepton final states at $\sqrt{s}=13$ TeV”. In: *Journal of High Energy Physics* 2021.208 (June 2021). DOI: [10.1007/JHEP07\(2021\)208](https://doi.org/10.1007/JHEP07(2021)208). URL: [https://doi.org/10.1007/JHEP07\(2021\)208](https://doi.org/10.1007/JHEP07(2021)208).
- [2] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 1–29. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2012.08.020>. URL: <https://www.sciencedirect.com/science/article/pii/S037026931200857X>.
- [3] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 30–61. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2012.08.021>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269312008581>.
- [4] P. Jackson and C. Rogan. “Recursive jigsaw reconstruction: HEP event analysis in the presence of kinematic and combinatoric ambiguities”. In: *Physical Review D* 96.11 (Dec. 2017). ISSN: 2470-0029. DOI: [10.1103/PhysRevD.96.112007](https://doi.org/10.1103/PhysRevD.96.112007). URL: <http://dx.doi.org/10.1103/PhysRevD.96.112007>.
- [5] A. Purnell. *Go on a particle quest at the first CERN webfest*. 2012. URL: <https://cds.cern.ch/record/1473657?ln=en> (visited on 09/01/2022).
- [6] C. Burgard. *Example: Standard Model of physics*. 2016. URL: <https://texample.net/tikz/examples/model-physics/> (visited on 09/01/2022).
- [7] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [8] M. Peskin and D. Schroeder. *An introduction to quantum field theory*. Boulder, CO: Westview, 1995. URL: <https://cds.cern.ch/record/257493>.
- [9] J. Hořejší. *Fundamentals of Electroweak Theory*. Charles University, Prague: The Karolinum Press, 2002.
- [10] N. Cabibbo. “Unitary Symmetry and Leptonic Decays”. In: *Phys. Rev. Lett.* 10 (12 June 1963), pp. 531–533. DOI: [10.1103/PhysRevLett.10.531](https://doi.org/10.1103/PhysRevLett.10.531). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.10.531>.

- [11] M. Kobayashi and T. Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. In: *Progress of Theoretical Physics* 49.2 (Feb. 1973), pp. 652–657. ISSN: 0033-068X. DOI: [10.1143/PTP.49.652](https://doi.org/10.1143/PTP.49.652). eprint: <https://academic.oup.com/ptp/article-pdf/49/2/652/5257692/49-2-652.pdf>. URL: <https://doi.org/10.1143/PTP.49.652>.
- [12] P. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [13] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.321>.
- [14] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. In: *Phys. Rev. Lett.* 13 (20 Nov. 1964), pp. 585–587. DOI: [10.1103/PhysRevLett.13.585](https://doi.org/10.1103/PhysRevLett.13.585). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.585>.
- [15] CMS Collaboration. *Summaries of CMS cross section measurements*. 2022. URL: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsCombined> (visited on 09/01/2022).
- [16] V. C. Rubin, Jr. Ford W. K., and N. Thonnard. “Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc).” In: 238 (June 1980), pp. 471–487. DOI: [10.1086/158003](https://doi.org/10.1086/158003).
- [17] A. Bosma. “21-cm line studies of spiral galaxies. I. Observations of the galaxies NGC 5033, 3198, 5055, 2841, and 7331.” In: 86 (Dec. 1981), pp. 1791–1824. DOI: [10.1086/113062](https://doi.org/10.1086/113062).
- [18] D. N. Spergel et al. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters”. In: *The Astrophysical Journal Supplement Series* 148.1 (Sept. 2003), pp. 175–194. DOI: [10.1086/377226](https://doi.org/10.1086/377226). URL: <https://doi.org/10.1086/377226>.
- [19] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1562>.
- [20] B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. In: *Sov. Phys. JETP* 7 (1958), pp. 172–173.
- [21] S. Martin. “A Supersymmetry Primer”. In: *Perspectives on Supersymmetry*. World Scientific, July 1998, pp. 1–98. DOI: [10.1142/9789812839657_0001](https://doi.org/10.1142/9789812839657_0001). URL: https://doi.org/10.1142/F9789812839657_0001.

- [22] H. Georgi and S. Glashow. “Unity of All Elementary-Particle Forces”. In: *Phys. Rev. Lett.* 32 (8 Feb. 1974), pp. 438–441. DOI: [10.1103/PhysRevLett.32.438](https://doi.org/10.1103/PhysRevLett.32.438). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.32.438>.
- [23] A. Leike. “The Phenomenology of extra neutral gauge bosons”. In: *Phys. Rept.* 317 (1999), p. 143. DOI: [10.1016/S0370-1573\(98\)00133-1](https://doi.org/10.1016/S0370-1573(98)00133-1). arXiv: [hep-ph/9805494](https://arxiv.org/abs/hep-ph/9805494).
- [24] H. Fritzsch and P. Minkowski. “Unified interactions of leptons and hadrons”. In: *Annals of Physics* 93.1 (1975), pp. 193–266. ISSN: 0003-4916. DOI: [https://doi.org/10.1016/0003-4916\(75\)90211-0](https://doi.org/10.1016/0003-4916(75)90211-0). URL: <https://www.sciencedirect.com/science/article/pii/0003491675902110>.
- [25] R. W. Robinett and Jonathan L. Rosner. “Mass scales in grand unified theories”. In: *Phys. Rev. D* 26 (9 Nov. 1982), pp. 2396–2419. DOI: [10.1103/PhysRevD.26.2396](https://doi.org/10.1103/PhysRevD.26.2396). URL: <https://link.aps.org/doi/10.1103/PhysRevD.26.2396>.
- [26] P. Langacker. “The physics of heavy Z' gauge bosons”. In: *Reviews of Modern Physics* 81.3 (Aug. 2009), pp. 1199–1228. DOI: [10.1103/revmodphys.81.1199](https://doi.org/10.1103/revmodphys.81.1199). URL: <https://doi.org/10.1103/revmodphys.81.1199>.
- [27] S. Dimopoulos and H. Georgi. “Softly broken supersymmetry and $SU(5)$ ”. In: *Nuclear Physics B* 193.1 (1981), pp. 150–162. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(81\)90522-8](https://doi.org/10.1016/0550-3213(81)90522-8). URL: <https://www.sciencedirect.com/science/article/pii/0550321381905228>.
- [28] J. Wess and B. Zumino. “Supergauge transformations in four dimensions”. In: *Nuclear Physics B* 70.1 (1974), pp. 39–50. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(74\)90355-1](https://doi.org/10.1016/0550-3213(74)90355-1). URL: <https://www.sciencedirect.com/science/article/pii/0550321374903551>.
- [29] T. Gherghetta, T. Kaeding, and G. Kane. “Supersymmetric contributions to the decay of an extra Z' boson”. In: *Physical Review D* 57.5 (Mar. 1998), pp. 3178–3181. DOI: [10.1103/PhysRevD.57.3178](https://doi.org/10.1103/PhysRevD.57.3178). URL: <https://doi.org/10.1103/PhysRevD.57.3178>.
- [30] M. Cvetič and P. Langacker. “ Z' Physics and Supersymmetry”. In: *Perspectives on Supersymmetry*. WORLD SCIENTIFIC, July 1998, pp. 312–331. DOI: [10.1142/9789812839657_0012](https://doi.org/10.1142/9789812839657_0012). URL: https://doi.org/10.1142/F9789812839657_0012.
- [31] M. Dugan, H. Georgi, and D. Kaplan. “Anatomy of a composite Higgs model”. In: *Nuclear Physics B* 254 (1985), pp. 299–326. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(85\)90221-4](https://doi.org/10.1016/0550-3213(85)90221-4). URL: <https://www.sciencedirect.com/science/article/pii/0550321385902214>.

- [32] Arkani-Hamed N., A. Dimopoulos, and G. Dvali. “The hierarchy problem and new dimensions at a millimeter”. In: *Physics Letters B* 429.3-4 (June 1998), pp. 263–272. DOI: [10.1016/S0370-2693\(98\)00466-3](https://doi.org/10.1016/S0370-2693(98)00466-3). URL: [https://doi.org/10.1016/S0370-2693\(98\)00466-3](https://doi.org/10.1016/S0370-2693(98)00466-3).
- [33] L. Randall and R. Sundrum. “Large Mass Hierarchy from a Small Extra Dimension”. In: *Phys. Rev. Lett.* 83 (17 Oct. 1999), pp. 3370–3373. DOI: [10.1103/PhysRevLett.83.3370](https://link.aps.org/doi/10.1103/PhysRevLett.83.3370). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.83.3370>.
- [34] CMS Collaboration. *CMS Beyond SM particles decaying 2(to) Higgs, top and Gauge bosons (B2G) Public Physics Results*. 2022. URL: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsB2G> (visited on 10/08/2022).
- [35] CMS Collaboration. *Summaries of CMS Exotica Results*. 2022. URL: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/SummaryPlotsEX013TeV> (visited on 10/08/2022).
- [36] CMS Collaboration. “Search for high mass dijet resonances with a new background prediction method in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of High Energy Physics* 2020 (5 May 2020), p. 33. DOI: [10.1007/JHEP05\(2020\)033](https://doi.org/10.1007/JHEP05(2020)033).
- [37] David J. Gross and R. Jackiw. “Effect of Anomalies on Quasi-Renormalizable Theories”. In: *Phys. Rev. D* 6 (2 July 1972), pp. 477–493. DOI: [10.1103/PhysRevD.6.477](https://link.aps.org/doi/10.1103/PhysRevD.6.477). URL: <https://link.aps.org/doi/10.1103/PhysRevD.6.477>.
- [38] B. Dobrescu. “Leptophobic Boson Signals with Leptons, Jets and Missing Energy”. In: (June 2015). arXiv: [1506.04435 \[hep-ph\]](https://arxiv.org/abs/1506.04435).
- [39] L. Evans and P. Bryant. “LHC Machine”. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08001–S08001. DOI: [10.1088/1748-0221/3/08/S08001](https://doi.org/10.1088/1748-0221/3/08/S08001). URL: <https://doi.org/10.1088/1748-0221/3/08/S08001>.
- [40] URL: <https://uspas.fnal.gov/materials/10MIT/Emittance.pdf>.
- [41] M. Benedikt et al. *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: [10.5170/CERN-2004-003-V-3](https://doi.org/10.5170/CERN-2004-003-V-3). URL: <http://cds.cern.ch/record/823808>.
- [42] E. Mobs. “The CERN accelerator complex - August 2018. Complexe des accélérateurs du CERN - Août 2018”. In: (Aug. 2018). General Photo. URL: <https://cds.cern.ch/record/2636343>.
- [43] “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *Journal of Instrumentation* 3 (2008). Also published by CERN Geneva in 2010, S08003. 437 p. DOI: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003). URL: <https://cds.cern.ch/record/1129811>.

- [44] S. Chatrchyan et al. “The CMS experiment at the CERN LHC”. In: *Journal of Instrumentation* 3 (2008), S08004. DOI: [10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004).
- [45] “The LHCb Detector at the LHC”. In: *Journal of Instrumentation* 3 (2008). Also published by CERN Geneva in 2010, S08005. DOI: [10.1088/1748-0221/3/08/S08005](https://doi.org/10.1088/1748-0221/3/08/S08005). URL: <https://cds.cern.ch/record/1129809>.
- [46] “The ALICE experiment at the CERN LHC”. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08002–S08002. DOI: [10.1088/1748-0221/3/08/s08002](https://doi.org/10.1088/1748-0221/3/08/s08002). URL: <https://doi.org/10.1088/1748-0221/3/08/s08002>.
- [47] URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults#Run_2_annual_charts_of_luminosit.
- [48] C. Fabjan and F. Gianotti. “Calorimetry for Particle Physics”. In: *Rev. Mod. Phys.* 75 (Oct. 2003), 1243–1286. 96 p. DOI: [10.1103/RevModPhys.75.1243](https://doi.org/10.1103/RevModPhys.75.1243). URL: <https://cds.cern.ch/record/692252>.
- [49] R. Wigmans. *Calorimetry: Energy Measurement in Particle Physics*. Oxford Science Publications, 2017. DOI: [0.1093/oso/9780198786351.001.0001](https://doi.org/0.1093/oso/9780198786351.001.0001).
- [50] URL: <https://wiki.physik.uzh.ch/cms/latex:tikz>.
- [51] W. Adam et al. “The CMS Phase-1 Pixel Detector Upgrade”. In: *Journal of Instrumentation* 16 (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].
- [52] “The CMS hadron calorimeter project: Technical Design Report”. In: (1997).
- [53] “The CMS high level trigger”. In: *Eur. Phys. J. C* 46 (2006), pp. 605–667. DOI: <https://doi.org/10.1140/epjc/s2006-02495-8>.
- [54] M. Brice, L. Vaillat, and L. Lazic. “Images of the CMS HCAL Barrel (HB)”. CMS Collection. Nov. 2008. URL: <https://cds.cern.ch/record/1431485>.
- [55] J. Mans et al. *CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter*. Tech. rep. Additional contact persons: Jeffrey Spalding, Fermilab, spalding@cern.ch, Didier Contardo, Universite Claude Bernard-Lyon I, contardo@cern.ch. Sept. 2012. URL: <http://cds.cern.ch/record/1481837>.
- [56] CMS Collaboration. “Identification and filtering of uncharacteristic noise in the CMS hadron calorimeter”. In: *Journal of Instrumentation* 5.03 (Mar. 2010), T03014–T03014. DOI: [10.1088/1748-0221/5/03/t03014](https://doi.org/10.1088/1748-0221/5/03/t03014). URL: <https://doi.org/10.1088/1748-0221/5/03/t03014>.
- [57] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: [10.23731/CYRM-2020-0010](https://doi.org/10.23731/CYRM-2020-0010). URL: <https://cds.cern.ch/record/2749422>.

- [58] F. Vasey et al. “The Versatile Link common project: feasibility report”. In: *Journal of Instrumentation* 7.01 (Jan. 2012), pp. C01075–C01075. doi: [10.1088/1748-0221/7/01/c01075](https://doi.org/10.1088/1748-0221/7/01/c01075). URL: <https://doi.org/10.1088/1748-0221/7/01/c01075>.
- [59] G. Cummings. “CMS HCAL VTRx-induced communication loss and mitigation”. In: *Journal of Instrumentation* 17.05 (May 2022), p. C05020. doi: [10.1088/1748-0221/17/05/c05020](https://doi.org/10.1088/1748-0221/17/05/c05020). URL: <https://doi.org/10.1088/1748-0221/17/05/c05020>.
- [60] O. Groettvik. “Installation, integration and first operating experiences of the ALICE ITS upgraded readout system”. In: *Journal of Instrumentation* 17.04 (Apr. 2022), p. C04003. doi: [10.1088/1748-0221/17/04/c04003](https://doi.org/10.1088/1748-0221/17/04/c04003). URL: <https://doi.org/10.1088/1748-0221/17/04/c04003>.
- [61] CMS Collaboration. “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”. In: *Journal of Instrumentation* 16.05 (May 2021), P05014. doi: [10.1088/1748-0221/16/05/p05014](https://doi.org/10.1088/1748-0221/16/05/p05014). URL: <https://doi.org/10.1088/1748-0221/16/05/p05014>.
- [62] A. M. Sirunyan et al. “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of Instrumentation* 13.10 (2018), P10005. doi: [10.1088/1748-0221/13/10/P10005](https://doi.org/10.1088/1748-0221/13/10/P10005). arXiv: [1809.02816](https://arxiv.org/abs/1809.02816) [hep-ex].
- [63] CMS Collaboration. “Particle-flow reconstruction and global event description with the CMS detector”. In: *Journal of Instrumentation* 12.10 (Oct. 2017), P10003–P10003. doi: [10.1088/1748-0221/12/10/p10003](https://doi.org/10.1088/1748-0221/12/10/p10003). URL: <https://doi.org/10.1088/1748-0221/12/10/p10003>.
- [64] The CMS Collaboration. “Description and performance of track and primary-vertex reconstruction with the CMS tracker”. In: *Journal of Instrumentation* 9.10 (Oct. 2014), P10009–P10009. doi: [10.1088/1748-0221/9/10/p10009](https://doi.org/10.1088/1748-0221/9/10/p10009). URL: <https://doi.org/10.1088/1748-0221/9/10/p10009>.
- [65] Pierre Billoir. “Progressive track recognition with a Kalman-like fitting procedure”. In: *Computer Physics Communications* 57.1 (1989), pp. 390–394. ISSN: 0010-4655. doi: [https://doi.org/10.1016/0010-4655\(89\)90249-X](https://doi.org/10.1016/0010-4655(89)90249-X). URL: <https://www.sciencedirect.com/science/article/pii/001046558990249X>.
- [66] CMS Collaboration. “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s}=13$ TeV”. In: *Journal of Instrumentation* 13.06 (June 2018), P06015–P06015. doi: [10.1088/1748-0221/13/06/p06015](https://doi.org/10.1088/1748-0221/13/06/p06015). URL: <https://doi.org/10.1088/1748-0221/13/06/p06015>.
- [67] W. Adam et al. “Reconstruction of electrons with the Gaussian-sum filter in the CMS tracker at the LHC”. In: *Journal of Physics G: Nuclear and Particle Physics* 31.9 (July 2005), N9–N20. doi: [10.1088/0954-3899/31/9/n01](https://doi.org/10.1088/0954-3899/31/9/n01). URL: <https://doi.org/10.1088/0954-3899/31/9/n01>.

- [68] M. Cacciari, G. Salam, and G. Soyez. “The anti- k_t jet clustering algorithm”. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 063–063. DOI: [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063). URL: <https://doi.org/10.1088/1126-6708/2008/04/063>.
- [69] S. Ellis and D. Soper. “Successive combination jet algorithm for hadron collisions”. In: *Physical Review D* 48.7 (Oct. 1993), pp. 3160–3166. DOI: [10.1103/physrevd.48.3160](https://doi.org/10.1103/physrevd.48.3160). URL: <https://doi.org/10.1103physrevd.48.3160>.
- [70] A. Sirunyan et al. “Pileup mitigation at CMS in 13 TeV data”. In: *Journal of Instrumentation* 15 (2020), P09018. DOI: [10.1088/1748-0221/15/09/p09018](https://doi.org/10.1088/1748-0221/15/09/p09018). arXiv: [2003.00503](https://arxiv.org/abs/2003.00503) [hep-ex].
- [71] D. Bertolini et al. “Pileup Per Particle Identification”. In: *Journal of High Energy Physics* 10 (2014), p. 059. DOI: [10.1007/JHEP10\(2014\)059](https://doi.org/10.1007/JHEP10(2014)059). arXiv: [1407.6013](https://arxiv.org/abs/1407.6013) [hep-ph].
- [72] A. Larkoski et al. “Soft drop”. In: *Journal of High Energy Physics* 2014.5 (May 2014). ISSN: 1029-8479. DOI: [10.1007/jhep05\(2014\)146](https://doi.org/10.1007/jhep05(2014)146). URL: [http://dx.doi.org/10.1007/JHEP05\(2014\)146](http://dx.doi.org/10.1007/JHEP05(2014)146).
- [73] “Jet energy scale and resolution performance with 13 TeV data collected by CMS in 2016-2018”. In: (2020). URL: <http://cds.cern.ch/record/2715872>.
- [74] Vardan Khachatryan et al. “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”. In: *Journal of Instrumentation* 12 (2017), P02014. DOI: [10.1088/1748-0221/12/02/P02014](https://doi.org/10.1088/1748-0221/12/02/P02014). arXiv: [1607.03663](https://arxiv.org/abs/1607.03663) [hep-ex].
- [75] JERC Subgroup CMS. *Jet Energy Resolution:Smearing Procedures*. 2022. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/JetResolution#Smearing_procedures (visited on 09/21/2022).
- [76] CMS Collaboration. “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector”. In: 14.07 (July 2019), P07004–P07004. DOI: [10.1088/1748-0221/14/07/p07004](https://doi.org/10.1088/1748-0221/14/07/p07004). URL: <https://doi.org/10.1088/1748-0221/14/07/p07004>.
- [77] C.G Lester and D.J Summers. “Measuring masses of semi-invisibly decaying particle pairs produced at hadron colliders”. In: *Physics Letters B* 463.1 (Sept. 1999), pp. 99–103. ISSN: 0370-2693. DOI: [10.1016/S0370-2693\(99\)00945-4](https://doi.org/10.1016/S0370-2693(99)00945-4). URL: [http://dx.doi.org/10.1016/S0370-2693\(99\)00945-4](http://dx.doi.org/10.1016/S0370-2693(99)00945-4).
- [78] M. Buckley et al. “Super-razor and searches for sleptons and charginos at the LHC”. In: *Physical Review D* 89.5 (Mar. 2014). ISSN: 1550-2368. DOI: [10.1103/physrevd.89.055020](https://doi.org/10.1103/physrevd.89.055020). URL: <http://dx.doi.org/10.1103/PhysRevD.89.055020>.
- [79] C. Rogan. URL: <http://restframes.com/> (visited on 09/05/2022).

- [80] EGamma POG. *Electron IDs*. 2020. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/EgammaRunIIRecommendations#Electrons_IDS (visited on 09/05/2022).
- [81] EGamma POG. *Electron Scale Factors*. 2020. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/EgammaRunIIRecommendations#Electron_Scale_Factors (visited on 09/05/2022).
- [82] EGamma POG. *Approved and ongoing electron and photon trigger scale factor measurements for Run2 data*. 2022. URL: <https://twiki.cern.ch/twiki/bin/viewauth/CMS/EgHLTScaleFactorMeasurements> (visited on 09/22/2022).
- [83] A.M. Sirunyan et al. “Performance of the reconstruction and identification of high-momentum muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of Instrumentation* 15.02 (Feb. 2020), P02027–P02027. ISSN: 1748-0221. DOI: [10.1088/1748-0221/15/02/p02027](https://doi.org/10.1088/1748-0221/15/02/p02027). URL: <http://dx.doi.org/10.1088/1748-0221/15/02/P02027>.
- [84] Muon POG. *Muon Legacy 2016*. 2021. URL: <https://twiki.cern.ch/twiki/bin/view/CMS/MuonLegacy2016> (visited on 09/05/2022).
- [85] Muon POG. *Muon Legacy 2017*. 2021. URL: https://twiki.cern.ch/twiki/bin/view/CMS/MuonLegacy2017#RECO_and_ID_efficiencies_AN1 (visited on 09/05/2022).
- [86] Muon POG. *Muon Legacy 2018*. 2021. URL: https://twiki.cern.ch/twiki/bin/view/CMS/MuonLegacy2018#Prompt_muons_AN1 (visited on 09/05/2022).
- [87] CMS Collaboration. “Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques”. In: *Journal of Instrumentation* 15.06 (June 2020), P06005–P06005. DOI: [10.1088/1748-0221/15/06/p06005](https://doi.org/10.1088/1748-0221/15/06/p06005). URL: <https://doi.org/10.1088/1748-0221/15/06/p06005>.
- [88] O. Colegrov, L. Gouskos, H. Huang, J. Incandela, J. Kieseler, Q. Li, M. Mozer, H. Qu, P. Sphicas, M. Stoye, M. Verzetti. *DeepAK8: Multi-class Boosted Jet Tagger*. 2018. URL: https://indico.cern.ch/event/695320/contributions/2850962/attachments/1632794/2603820/status_deepAK8_CMSHFT_2018-04-13_hqu.pdf (visited on 09/08/2022).
- [89] JME POG. *MET Corrections and Uncertainties for Run-II*. 2021. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/MissingETRun2Corrections#xy_Shift_Correction_MET_phi_modu.
- [90] A. M. Sirunyan et al. “Search for ZZ resonances in the $2\ell 2\nu$ final state in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of High Energy Physics* 2018.3 (2018). DOI: [https://doi.org/10.1007/JHEP03\(2018\)003](https://doi.org/10.1007/JHEP03(2018)003).
- [91] The SciPy Community. *scipy.interpolate.RBFInterpolator*. 2022. URL: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.RBFInterpolator.html> (visited on 10/10/2022).

- [92] ATLAS Collaboration. “Search for new resonances in mass distributions of jet pairs using 139 fb^{−1} of pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector”. In: *Journal of High Energy Physics* 2020 (3 Feb. 2020), p. 145. URL: [https://doi.org/10.1007/JHEP03\(2020\)145](https://doi.org/10.1007/JHEP03(2020)145).
- [93] “Search for narrow resonances in dijet final states at $\sqrt{s}$ = 8 TeV with the novel CMS technique of data scouting”. In: *Phys. Rev. Lett.* 117 (3 July 2016), p. 031802. DOI: [10.1103/PhysRevLett.117.031802](https://doi.org/10.1103/PhysRevLett.117.031802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.031802>.
- [94] “Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s}$ = 13TeV”. In: *Phys. Rev. D* 100 (11 Dec. 12), p. 112007. DOI: [10.1103/PhysRevD.100.112007](https://doi.org/10.1103/PhysRevD.100.112007). URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.112007>.
- [95] CMS Collaboration. “The performance of the CMS muon detector in proton-proton collisions at $\sqrt{s}$ = 7 TeV at the LHC”. In: *Journal of Instrumentation* 8.11 (Nov. 2013), P11002–P11002. DOI: [10.1088/1748-0221/8/11/p11002](https://doi.org/10.1088/1748-0221/8/11/p11002). URL: <https://doi.org/10.1088/1748-0221/8/11/p11002>.
- [96] CMS Collaboration. “Search for heavy resonances decaying to a pair of Lorentz-boosted Higgs bosons in final states with leptons and a bottom quark pair at $\sqrt{s}$ = 13 TeV”. In: *Journal of High Energy Physics* 2022.5 (May 2022). DOI: [10.1007/JHEP05\(2022\)005](https://doi.org/10.1007/JHEP05(2022)005). URL: [https://doi.org/10.1007/JHEP05\(2022\)005](https://doi.org/10.1007/JHEP05(2022)005).

Appendices

Appendix A

Scale Factors

Scale factors re-weight Monte Carlo simulation to match the response observed in data. They are generally derived to account for reconstruction, identification, trigger, and btagging inefficiencies. The scale factor itself is the ratio of the data efficiency, binned in p_T and η , to the Monte Carlo efficiency. The efficiencies are generically measured using the Tag and Probe method [95].

A.1 Muon Scale Factors

Prior to this analysis, muon trigger scale factors did not exist for the trigger and ID combination. We derived the muon scale factors for this analysis.

The p_T and η bins used in the efficiency measurement are given in Table A.1. The bins below 200 GeV in p_T and all η bins align with those used in the official muon trigger scale factor production for 2017 and 2018. The additional bin from 200 GeV to 500 GeV was added to capture any tail effects due to our boosted signal. Figure A.1 shows the expected p_T of the leading muon and subleading muon in the Z candidate distribution for the Leptophobic Z' signal MC. The efficiency measurement was not pushed to higher momenta, and the binning was kept coarse, due to the limited statistics of both the data and the MC in the high p_T region. Furthermore, such efforts are not necessary when the yields in the high p_T region are negligible at our expected sensitivity.


Table A.2 summarizes the tag and probe requirements for the efficiency measurements. The probe muon's p_T requirement matches the subleading muon p_T requirement in the Z candidate reconstruction (Section 6.4.3). The calculation of the scale factor extends this low to account for the contribution of both muons in the Z candidate to the total efficiency.

The muon ID scale factors are derived centrally by CMS. Scale factors are applied for both muons in the Z candidate, with their product serving as the event weight. Figure A.3 shows the muon ID scale factors. The 2016 scale factors were derived for two separate data-taking periods, and must be applied as a luminosity-weighted sum.

Bins for efficiency measurements

p_T (GeV)	$ \eta $
[20.0, 25.0)	[0.0, 0.9)
[25.0, 30.0)	[0.9, 1.2)
[30.0, 40.0)	[1.2, 2.1)
[40.0, 50.0)	[2.1, 2.4)
[50.0, 55.0)	
[55.0, 60.0)	
[60.0, 120.0)	
[120.0, 200.0)	
[200.0, 500.)	

Table A.1: Binning used in the efficiency measurements and scale factor generation.

Figure A.1: Leading and subleading muon p_T in the signal region for a signal cross section of 10 fb. Boosted signal is possible, hence the addition of the last bin. No trigger scale factors applied.

Tag and Probe Selections

Tag	Probe
Tight ID	Tight ID
$p_T > 20$ GeV	$p_T > 20$ GeV
$ \eta < 2.4$	$ \eta < 2.4$
	HLT_Mu50_v*
	OR HLT_TkMu50_v* (2016)
	OR HLT_TkMu100_v* (2017 and 2018)

Table A.2: Summary of the tag and probe selections.




Figure A.2: Privately generated muon trigger scale factors for 2016 (top), 2017 (middle), and 2018 (bottom).




Figure A.3: Muon ID scale factors applied. The top figure shows the luminosity weighted 2016 scale factors [84]. The middle shows 2017 [85], and the bottom shows 2018 [86].

A.2 Electron Scale Factors

The electron reconstruction scale factors are provided centrally by CMS. The reconstruction scale factors are applied based on the selected electron's p_T and η , and are shown in Fig. A.4. The electron ID scale factors are also provided centrally and are applied with the same criteria as the reconstruction scale factors, and are shown in Fig. A.5. In either case, if the electron's p_T is greater than the final p_T bin, the scale factor from the highest p_T bin in the appropriate η bin is applied.

CMS does not produce electron trigger scale factors, but does maintain lists of approved trigger scale factors developed for other analyses [82]. We apply the HEEP V70 ID scale factors that correspond with our trigger menu, generated by Jae Sung Kim. Since the HEEP V70 ID is a tighter requirement than our own, this scale factor application is an overestimate of the true inefficiency. This translates to a larger systematic impact, though still within the statistical uncertainty on the $t\bar{t}$ normalization.

A.3 Btagging Scale Factors

The btagging scale factors are the same as those in Ref. [96], received via private correspondence. Two types of btagging SFs exist: $t\bar{t}$ sample-specific weights to account for top p_T reweighting, and a generic set for all other processes. They are not derived via tag and probe, but are instead derived following the procedure in Ref. [87]. Figure A.6 shows the applied btagging scale factors. These weights only depend on the jet p_T .




Figure A.4: Electron reconstruction scale factors applied. From top to bottom: 2016, 2017, 2018. [81].




Figure A.5: Electron ID scale factors applied. From top to bottom: 2016, 2017, 2018. [80].




Figure A.6: Left column: Generic btagging scale factors. Right column: Btagging scale factors for $t\bar{t}$ MC.

Appendix B

Control Plots for the electron and muon data set used to estimate $t\bar{t}$ - Preselection and Jet Mass Signal Region

These plots are all generated using the fake Z to electron and muon control region. Plots are derived combining the SingleMuon and SingleElectron/EGamma data sets and the MC in the fake Z channel, as described in Section 7.1. Plots are presented at preselection level, in the jet mass signal region. These plots represent the distributions that would be scaled by the $e\mu$ ratio to serve as the estimate of the $t\bar{t}$ background in the analysis' signal region.




Figure B.1: Z to $e\mu$ at preselection in the jet mass signal region featuring H candidate kinematic distributions.




Figure B.2: Z to $e\mu$ at preselection in the jet mass signal region featuring dilepton fake Z candidate kinematic distributions.




Figure B.3: Z to $e\mu$ at preselection in the jet mass signal region featuring the RJR mass estimators.




Figure B.4: Z to $e\mu$ at preselection in the jet mass signal region featuring mass variables, and btagging values.




Figure B.5: Z to $e\mu$ at preselection in the jet mass signal region kinematic for the leading lepton.




Figure B.6: Z to $e\mu$ at preselection in the jet mass signal region kinematic for the subleading lepton.




Figure B.7: Z to $e\mu$ at preselection in the jet mass signal region kinematics for the muons selected to serve as the background estimation for the $\mu\mu$ channel.




Figure B.8: Z to $e\mu$ at preselection in the jet mass signal region kinematics for the electrons selected to serve as the background estimation for the $\mu\mu$ channel.




Figure B.9: Z to $e\mu$ at preselection in the jet mass signal region ΔR plots for Higgs candidate jet and Z constituents.




Figure B.10: Z to $e\mu$ at preselection in the jet mass signal region $\Delta\phi$ plots for boson candidates and MET.

Appendix C

α -method validation region closure plots

Additional closure plots for the validation region. Good data/MC agreement is observed, lending credence to the use of fits to Monte Carlo.




Figure C.1: Additional closure plots in the validation region. Same selection as preselection signal region except the p_T^{miss} cut is reversed.

Appendix D

Impact of low background on Signal Injection Tests

The negative bias observed in the signal injection test is due to the low expected number of background events. When performing the signal injection tests, the measured signal strength is allowed to take negative values. If the background yield is low, the signal strength can acquire negative values to describe the generated pseudoexperiments, which can fluctuate within the provided uncertainties.

This behavior was explored in a limited version of the statistical model by scaling the expected background yields. With the background yield artificially inflated, the bias disappears. This verifies that the bias is only due to the low yield. In a true significance extraction, the signal strength is forbidden to take negative values.

The signal model considered is 8 bins with 200 GeV width in the Z' RJR mass estimator spectrum between 1400 and 3000 GeV. This range was chosen to avoid the turn-on curve of the backgrounds, and to remain in the region where the background is understood. The methods are the same as those considered in Section 9.3.3.

Figures D.1 through D.3 show pairs of a nominal signal injection pull and its background-scaled counterpart. The nominal distributions are considered biased, and the scaled versions are bias free. The pulls from the scaled distributions are now centered at 0, as expected for an unbiased model.




Figure D.1: Comparison between the unscaled (left) and scaled by 10 (right) pull distributions for the $m_{Z'} = 2000$, $m_{N_D} = 400$, $m_{N_S} = 200$ signal mass point at the median expected limit injected signal strength.




Figure D.2: Comparison between the unscaled (left) and scaled by 10 (right) pull distributions for the $m_{Z'} = 4000$, $m_{N_D} = 800$, $m_{N_S} = 200$ signal mass point at the median expected limit injected signal strength. Slight excess at negative pull values gone in the scaled version.




Figure D.3: Comparison between the unscaled (left) and scaled by 10 (right) pull distributions for the $m_{Z'} = 5500$, $m_{N_D} = 1800$, $m_{N_S} = 200$ signal mass point at 0.0 injected signal strength. Slight excess at negative pull values gone in the scaled version.