

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Measurements of the Higgs Boson through Vector Boson Scattering and Software and Computing for Exascale Data Science

Permalink

<https://escholarship.org/uc/item/9pm173nr>

Author

Guiang, Jonathan Kasuke

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Measurements of the Higgs Boson through Vector Boson Scattering and Software and
Computing for Exascale Data Science

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Physics

by

Jonathan Kasuke Guiang

Committee in charge:

Professor Frank Würthwein, Chair
Professor R. Sekhar Chivukula
Professor Albert Hsiao
Professor Arun Kumar
Professor Michael Norman
Professor Avraham Yagil

2024

Copyright

Jonathan Kasuke Guiang, 2024

All rights reserved.

The Dissertation of Jonathan Kasuke Guiang is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

This dissertation is dedicated to my wife Rosemary and the future we are building together; to my parents, Orlando and Cynthia Guiang, and my sister, Christina Guiang; to grandpa and grandma, John and Tsuneko Hashiguchi; to abuelo and abuela, Rufino and Aurelia Guiang; to the seas crossed, the hardship, the sacrifices, but also the love, the food, the adventures, and whatever comes next.

I love you all very much.

EPIGRAPH

*To see a World in a Grain of Sand
And a Heaven in a Wild Flower
Hold Infinity in the palm of your hand
And Eternity in an hour*
(William Blake, *Auguries of Innocence*, 1863)

In the particular is contained the universal.
(James Joyce, interview with Arthur Power, 1921)

*Talent is insignificant.
I know a lot of talented ruins.
Beyond talent lie all
the usual words: discipline, love, luck,
but, most of all, endurance.*
(James Baldwin, *The Paris Review*, 1984)

TABLE OF CONTENTS

Dissertation Approval Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures	ix
List of Tables	xiv
Acknowledgements	xv
Vita	xix
Abstract of the Dissertation	xx
Introduction	1
Chapter 1 A Beautiful Theory	6
1.1 The Standard Model	6
1.1.1 Feynman diagrams	7
1.1.2 Cross sections	9
1.1.3 The Higgs boson	12
1.2 Selected open questions	14
1.2.1 What is dark matter?	14
1.2.2 What is dark energy?	15
1.2.3 What is the true nature of the Higgs boson?	15
Chapter 2 A Grand Apparatus	19
2.1 The Large Hadron Collider	19
2.1.1 Collision energy	20
2.1.2 Luminosity and pileup	21
2.2 The Compact Muon Solenoid	24
2.2.1 Overview	25
2.2.2 Superconducting solenoid	26
2.2.3 Silicon tracker	28
2.2.4 Electromagnetic calorimeter (ECAL)	30
2.2.5 Hadronic calorimeter (HCAL)	31
2.2.6 Muon chambers	32
2.2.7 Trigger system	34
2.3 The high luminosity era	35
2.4 Acknowledgements	36

Chapter 3	Studying the Higgs boson through vector boson scattering at CMS	37
3.1	New physics and the Higgs boson	37
3.1.1	The κ -framework	38
3.2	Producing Higgs bosons with vector boson scattering	39
3.3	Choosing a final state	40
3.4	Reconstruction	41
3.4.1	Jets	42
3.4.2	Leptons	45
3.4.3	Missing transverse momentum	46
3.5	Identifying backgrounds	46
3.6	Simulation	47
3.7	Event selection	50
3.8	Background estimation	51
3.9	Systematic uncertainties	53
3.10	Statistical interpretation	54
3.10.1	Probability and likelihood	54
3.10.2	Maximum likelihood estimate	55
Chapter 4	Determination of λ_{WZ} through VBS WH production	59
4.1	Looking for a sign	59
4.2	The signal	60
4.2.1	Signal characteristics	60
4.3	The backgrounds	62
4.4	Event selection	63
4.4.1	Triggers and preselection	63
4.4.2	Signal region	64
4.5	Background estimation	65
4.6	Results	69
4.7	Acknowledgements	70
Chapter 5	Measurement of κ_{VV} through VBS VVH production	73
5.1	Chasing couplings	73
5.2	The signal	73
5.2.1	Signal characteristics	74
5.3	The backgrounds	75
5.3.1	QCD PARTICLENET resampling	75
5.4	Event selection	78
5.4.1	Triggers and preselection	78
5.4.2	ABCDNET	80
5.4.3	Signal region	82
5.5	Background estimation	84
5.6	Results	87
5.7	Next steps	88
5.8	Acknowledgements	89

Chapter 6	Line Segment Tracking	91
6.1	The HL-LHC computing challenge	91
6.1.1	Track reconstruction	92
6.2	The line segment tracking algorithm	93
6.3	Improving LST with machine learning	95
6.3.1	Training	95
6.3.2	Results	98
6.4	Next steps	98
6.5	Acknowledgements	100
Chapter 7	Exascale Cyberinfrastructure	101
7.1	The HL-LHC data challenge	101
7.2	Software defined networking	102
7.3	Rucio	103
7.4	SENSE	103
7.5	Rucio-SENSE interoperation model	104
7.6	Demonstrations	105
7.7	Next steps	107
7.8	Acknowledgements	108
Appendix A	CMS coordinate system	109
Bibliography	111

LIST OF FIGURES

Figure 1.1.	The Standard Model of particle physics.	8
Figure 1.2.	The fundamental vertices in the Standard Model.	9
Figure 1.3.	Different versions of the QED vertex.	10
Figure 1.4.	Feynman diagram for beta decay, which is the mechanism behind the radioactive decay of carbon-14 (C-14), which enables carbon dating.	10
Figure 1.5.	Leading-order Feynman diagram for electron-electron scattering.	12
Figure 1.6.	The totality of the cross-section measurements performed with data from the CMS Experiment.	17
Figure 1.7.	A selection of cross-section measurements performed with data from the ATLAS Experiment.	18
Figure 2.1.	A diagram illustrating the depth of the LHC beneath the surface and an aerial photograph of the entire CERN complex.	20
Figure 2.2.	The CERN accelerator complex illustrated in detail.	21
Figure 2.3.	A cutaway diagram illustrating the internals of a dipole magnet inside the LHC tunnel and a photograph of a decommissioned dipole magnet with a physicist for scale.	22
Figure 2.4.	Distribution of the average number of interactions per bunch crossing within the CMS detector.	23
Figure 2.5.	A detailed cutaway diagram of CMS, with each subdetector labeled.	24
Figure 2.6.	CMS with one of the endcaps separated from the main body of the experiment, pictured from the top and side.	25
Figure 2.7.	CMS in the closed configuration, pictured from the top and side.	25
Figure 2.8.	The field lines of an ideal solenoid and those of the CMS magnet.	27
Figure 2.9.	A photograph of a section of the CMS magnet and a diagram of the cross-section of a Rutherford cable.	28
Figure 2.10.	A photograph of the CMS silicon tracker inner forward pixel modules and the outer barrel layers with a physicist in the background for scale.	29
Figure 2.11.	A detailed cutaway diagram of the CMS Phase-1 tracker.	30

Figure 2.12.	A photograph of the CMS electromagnetic calorimeter (ECAL) after half of the modules had been installed, and a physicist posed with the lead tungstate scintillator crystals used in the ECAL.	31
Figure 2.13.	A photograph of the CMS hadronic calorimeter (HCAL) and the World War II artillery shell casings that were melted down to supply the brass. ...	32
Figure 2.14.	The cross-section of CMS in the r - z and r - ϕ planes, respectively.	33
Figure 2.15.	A photograph of a section of the barrel muon chambers and a diagram of a drift tube muon detector.	34
Figure 2.16.	A collision event with high pileup for Run 2 versus another with HL-LHC-like pileup, both recorded in 2016.	35
Figure 3.1.	The leading-order Feynman diagram for VBS production of a Higgs boson and two vector bosons.	39
Figure 3.2.	A transverse slice of CMS, showing the signals left by different kinds of particles in each subdetector.	41
Figure 3.3.	Illustration of two quarks hadronizing into two distinct jets.	42
Figure 3.4.	A sketch of the hadronization and identification of a b quark.	43
Figure 3.5.	Illustration of a generic boson with a small, moderate, and large Lorentz boost, respectively, decaying to two jets.	44
Figure 3.6.	VBS jets signature.	45
Figure 3.7.	Event display for a simulated VBS WH signal and $t\bar{t}$ event.	47
Figure 3.8.	Detailed diagram of the generation of a single $pp \rightarrow t\bar{t} + H$ event.	48
Figure 3.9.	Diagram of the CMS collision data and simulation processing workflow. .	49
Figure 3.10.	Diagram of the evolution of CMS data from collection to post-processing into the data used in analysis.	51
Figure 3.11.	Diagram of the data processing workflow for a typical measurement with CMS data.	52
Figure 3.12.	A generic ABCD configuration.	53
Figure 3.13.	The probability and likelihood.	55
Figure 3.14.	The likelihood with the Gaussian approximation of the likelihood overlaid.	57

Figure 3.15.	The likelihood and confidence interval plotted for the “signal exists” and “signal does not exist” scenarios.	58
Figure 4.1.	Leading-order Feynman diagrams for VBS WH production.	60
Figure 4.2.	Generator-level S_T distribution.	62
Figure 4.3.	The signal cross-section plotted as a function of enhancements to κ_W	63
Figure 4.4.	Event display for a simulated VBS WH event.	64
Figure 4.5.	Feynman diagrams the leading ($t\bar{t}$) and subleading (W+jets and single top quark) backgrounds.	65
Figure 4.6.	The M_{jj} and $\Delta\eta_{jj}$ distributions for the VBS jets.	66
Figure 4.7.	The S_T distribution.	67
Figure 4.8.	The $H \rightarrow b\bar{b}$ fat jet candidate soft drop mass and PARTICLENET $X \rightarrow b\bar{b}$ score distributions.	68
Figure 4.9.	A sketch of regions A, B, C, and D used in the background estimation method.	69
Figure 4.10.	A 2-dimensional histogram of the ABCD variables M_{SD} and $\Delta\eta_{jj}$ with a profile of M_{SD}	70
Figure 4.11.	The data and MC yields plotted in regions A, B, C, and D.	71
Figure 4.12.	The 1D and 2D exclusion significance for BSM λ_{WZ} scenarios.	72
Figure 5.1.	Leading-order Feynman diagrams for VBS VVH in the all-hadronic final state.	74
Figure 5.2.	The M_{jj} and $\Delta\eta_{jj}$ distribution for the VBS jets.	75
Figure 5.3.	The p_T distribution for each of the three VBS VVH AK8 jets.	76
Figure 5.4.	The PARTICLENET regressed mass distribution for each of the three VBS VVH AK8 jets.	76
Figure 5.5.	Feynman diagrams for QCD multijet and $t\bar{t}$ production.	77
Figure 5.6.	The PARTICLENET score distribution for each of the three VBS VVH AK8 jets plotted in data and MC before QCD resampling.	77
Figure 5.7.	The approximation of the PARTICLENET probability density functions. . .	79

Figure 5.8.	The PARTICLENET score distribution for each of the three VBS VVH AK8 jets plotted in data and MC after QCD resampling.	80
Figure 5.9.	The input features for ABCDNET plotted as histograms normalized to unity.	83
Figure 5.10.	The average loss at each epoch and ROC curve for the selected ABCDNET model.	84
Figure 5.11.	A two-dimensional histogram binned in ABCDNET score and $\Delta\eta_{jj}$	85
Figure 5.12.	The ABCDNET score plotted in data and MC at the Preselection.	86
Figure 5.13.	A cartoon of the ABCD regions.	87
Figure 5.14.	The ABCD regions used for closure tests.	88
Figure 5.15.	Regions A, B, C, and D used in the ABCD method.	89
Figure 5.16.	The 95% confidence level limit plotted as a function of κ_{VV} for the all-hadronic channel (left) and all channels combined (right).	90
Figure 6.1.	HL-LHC CPU usage projections at CMS.	91
Figure 6.2.	Simple illustration of particles leaving hits in a multi-layered tracker.	92
Figure 6.3.	An illustration of the CMS Phase-II silicon tracker with a diagram of a p_T -module.	94
Figure 6.4.	LST track-finding efficiency and fake rate plotted as a function of η	96
Figure 6.5.	The components of a T5 and the DNN architecture.	97
Figure 6.6.	The average loss after each epoch and ROC curve for the DNN.	97
Figure 6.7.	LST efficiency for all TCs plotted as a function of p_T and η	98
Figure 6.8.	LST fake rate for all TCs plotted as a function of p_T and η	99
Figure 6.9.	LST efficiency for all TCs plotted as a function of r_{vertex}	99
Figure 6.10.	LST throughput plotted for different numbers of parallel CUDA streams. .	100
Figure 7.1.	HL-LHC tape usage projections at CMS.	102
Figure 7.2.	Rucio-SENSE site configuration.	104
Figure 7.3.	Rucio-SENSE interoperation workflow.	106

Figure 7.4. Rucio-SENSE interoperation throughput with 2 sites in the testbed. 107

Figure 7.5. Rucio-SENSE interoperation throughput with 3 sites in the testbed. 108

Figure A.1. The CMS coordinate system. 110

LIST OF TABLES

Table 4.1.	VBS WH signal region yields.	71
Table 5.1.	ABCDNET hyperparameters.	82
Table 5.2.	VBS VVH ABCD yields.	85
Table 5.3.	VBS VVH ABCD cross-check yields.	90
Table 6.1.	LST track-finding steps and object definitions.	95

ACKNOWLEDGEMENTS

Recorded in this dissertation, in great detail, are my contributions to a grand scientific enterprise. What is not documented, however, are the struggles, doubts, heartache, frustrations, and so much more that lie beyond the histograms, tables, and physics prose. Through each of these challenges, I am fortunate to have been surrounded by incredible people who ran beside me in the best of times, and dragged me along in the worst. Please find below my best attempt at thanking them, in mere sentences, for making all of this—and all that is to come—possible.

Foremost, I owe boundless thanks to my incredible family. First, to Mom and Dad, thank you for raising me in a home full of love and imagination and for all of your work and sacrifice upon which my life, career, and now this dissertation have been built upon. And to Christina, my big sister: thank you for your strength and courage, for all of our memories growing up together (except, perhaps, those that ended with my injury), and for inspiring me to be bold. I am also fortunate to have spent a great deal of time with my grandparents, John and Tsuneko Hashiguchi—thank you, grandma and grandpa, for teaching me to be good and honest and for showing me how to persevere. Amidst all of my “found” family—family by bond, not blood—I must, most of all, thank my uncle, John Sanders: thank you for putting up with your briny nickname, for every phone call, for your humor (even when it is at my expense), and for helping me at every stage of my career. Finally, to my wife Rosemary, who has brightened my life in every way: thank you for your kindness and compassion, for everything that you are, I love you more than anything. I hope that today and on every day after, I can make you all proud.

I am equally grateful for the many mentors I have had in my relatively short academic career, organized here by the institution at which we met. Starting with UCSD, I thank Prof. Frank Würthwein, my advisor, who believed in me, guided me through all of my achievements, and flew me around the world—thank you, Frank, for everything. And to Prof. Avi Yagil, who served as my co-advisor: thank you for challenging, uplifting, and educating me. I also owe many thanks to Prof. Philip Chang who, in many ways, was also my co-advisor. For hours of conversation via Zoom and over 130 000 Skype messages, for even more hours of physics and

C++ lessons, and for your nearly endless patience: thank you, Philip. Next, from my alma mater, UCSB, I am ever indebted to Prof. Claudio Campagnari for taking me in and completely altering the course of my life. Claudio gave me purpose during what was, probably unbeknownst to him, a very difficult time for me personally—thank you, Claudio; without you, this dissertation simply would not exist. Also, Claudio’s former students and postdocs were my first and most formative mentors. Thanks, in particular, to Dr. Nick Amin for teaching me everything I know. Thanks also Dr. Bennett Marsh, who mentored me across multiple projects—I am a better leader, person, and scientist thanks to Nick and Bennett. And to Prof. Indara Suarez, who welcomed me into particle physics and pushed me to meet my true potential: thank you for the high-octane tour of Boston in your Mustang, your thoughtful mentorship, and your constant support.

I must also thank all of the members, past and present, of the illustrious “Surf n’ Turf” (SNT) empire, whose dominion stretches far across the western world. It has been a privilege to serve as your Data Librarian, and I am deeply grateful to each of you, though, unfortunately, I do not have space to list all of you. Therefore, my thanks are extended, but not limited to the following SNT’ers: Prof. Frank Golf, Dr. Daniel Spitzbart, Dr. Sam May, Aashay Arora, Dr. Slava Krutelyov, and former SNT’er Jerry Ling.

In addition to all of the aforementioned physicists, I have had the privilege of working alongside a number of incredible engineers whom I would like to mention here. In particular, I would like to thank Diego Dávila for patiently teaching some computer science to a budding physicist and for being a great mentor and friend. Thanks also to Terrence Martin for keeping our equipment running, teaching me how to deadlift, and for your amazing board game parties!

With the due date of this dissertation rapidly approaching, I would like to end this section by acknowledging some of my incredible friends. To Joey Incandela, my physics brother (and best man): thank you for the many nights of gaming (despite the 3-hour time difference), for keeping me sane, and for your compassion and spirit. To Gabe Hernandez, my ex-physics brother (and ex-roommate): thank you for the silly inside jokes, for your constant guidance and encouragement, and for accompanying me to the base of Mt. San Jacinto, where I finished

writing this dissertation amidst pizza, cookies, and the pines. To Grady Kestler, last but not least, my partner in crime: thank you for Sunday Night Games, for getting me through Stat. Mech., and for your camaraderie and insight. And to everyone else who has helped me along my journey: please know that I am sincerely thankful, and that I will endeavor to honor all of you with what I can contribute to the world.

Through the Sloan Scholar program at UCSD, the Alfred P. Sloan Foundation provided financial and moral support for the work presented in this dissertation. To Shana Slebioda, Sinai Cota, and the Sloan Foundation: thank you!

Chapter 2 describes the Large Hadron Collider (LHC) and the CMS Experiment. The figures used in this chapter are materials produced by or for CERN and were released by CERN for informational use under CERN copyright. The descriptions of the LHC and CMS were derived from publicly available material and the relevant technical design reports.

Chapter 4 describes a search for the production of a W and Higgs boson via vector boson scattering, using data recorded by the CMS experiment from 2016 to 2018. It is a partial reproduction of the paper “Study of WH production through vector boson scattering and extraction of the relative sign of the W and Z couplings to the Higgs boson in proton-proton collisions at $\sqrt{s} = 13$ TeV,” submitted to PLB (arXiv:2405.16566). My individual contribution to this work was significant. I participated in its ideation by working with Prof. Philip Chang (without telling my advisor, at first) to make the initial calculations of the feasibility of the measurement. Then, I personally implemented the analysis that measures the relative sign of the HWW and HZZ couplings, from basic selections, to the optimization of the signal region, to the background estimation and evaluation of systematic uncertainties.

Chapter 5 describes a search for the production of a Higgs boson and two vector bosons via vector boson scattering, using data recorded by the CMS experiment from 2016 to 2018. It is a partial reproduction of a paper being prepared for submission for publication. I worked as part of a multi-institutional team that, together, implemented the three main channels of this analysis. However, I was solely responsible for the channel described in Chapter 5. On top of

implementing the analysis itself (from selections to systematic uncertainties), I also implemented the training pipeline for the deep neural network used for the final result in both my work and the other most sensitive channel of the combined analysis.

Chapter 6 describes improvements to the Line Segment Tracking algorithm using machine learning. It is a partial reproduction of the paper “Improving tracking algorithms with machine learning: a case for line-segment tracking at the High Luminosity LHC,” in the proceedings of Connecting the Dots 2023 (arXiv:2403.13166). This work is part of a larger effort to produce a highly parallelizable line-segment based tracking (LST) algorithm. My personal contribution included the development and training of a machine learning algorithm towards improving track-finding performance, and the eventual integration of it into the LST algorithm.

Chapter 7 describes the development of software-defined networking (SDN) capabilities within the LHC data distribution cyberinfrastructure. It is a partial reproduction of the paper “Managed Network Services for Exascale Data Movement Across Large Global Scientific Collaborations,” in the proceedings of XLOOP 2022 (doi:10.1109/XLOOP56614.2022.00008). I personally developed the first version of the software that implements the interoperation of SDN into the LHC data management software, with guidance from research engineers at the SDSC and ESNet. I also configured the testbed of this system and orchestrated some of the first proof-of-concept tests. Diego Dávila, Aashay Arora, Tom Lehman, and Xi Yang were all instrumental in the success and continued progress of this work.

VITA

2019 Bachelor of Science, University of California Santa Barbara
2019–2020 Teaching Assistant, University of California San Diego
2020–2024 Research Assistant, University of California San Diego
2023 Master of Science, University of California San Diego
2024 Doctor of Philosophy, University of California San Diego

PUBLICATIONS

CMS Collaboration (2024). “Study of WH production through vector boson scattering and extraction of the relative sign of the W and Z couplings to the Higgs boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”. Submitted to PLB. arXiv:2405.16566.

Guiang, J. et al. (2023). “Improving tracking algorithms with machine learning: a case for line-segment tracking at the High Luminosity LHC”. In 8th International Connecting The Dots Workshop. arXiv:2403.13166.

Aashay Arora, Jonathan Guiang, Diego Davila, Frank Würthwein, Justas Balcas, & Harvey Newman (2023). “400Gbps benchmark of XRootD HTTP-TPC”. In 26th International Conference on Computing in High Energy & Nuclear Physics. EPJ Web of Conferences.

Würthwein, F., Guiang, J. et al. (2022). “Managed Network Services for Exascale Data Movement Across Large Global Scientific Collaborations”. In 2022 4th Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing. IEEE Computer Society.

Guiang, J. et al. (2022). “Integrating End-to-End Exascale SDN into the LHC Data Distribution Cyberinfrastructure”. In Practice and Experience in Advanced Research Computing. Association for Computing Machinery.

Fajardo, E., ..., Guiang, J. et al. (2020). “Moving the California distributed CMS XCache from bare metal into containers using Kubernetes”. In 24th International Conference on Computing in High Energy & Nuclear Physics. EPJ Web of Conferences.

Ball, A., ..., Guiang, J. et al. (2020). “Search for millicharged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV”. Physical Review D, 102(3).

ABSTRACT OF THE DISSERTATION

Measurements of the Higgs Boson through Vector Boson Scattering and Software and Computing for Exascale Data Science

by

Jonathan Kasuke Guiang

Doctor of Philosophy in Physics

University of California San Diego, 2024

Professor Frank Würthwein, Chair

This dissertation presents the analyses of WH and VVH production through vector boson scattering (VBS). The VBS WH analysis excludes scenarios where the HWW and HZZ couplings have opposite signs beyond 5 standard deviations. The VBS VVH analysis places limits on the HHVV coupling between -0.03 and 2.04 times the Standard Model value. Both analyses are based on proton-proton collision data recorded by the CMS experiment at the CERN LHC from 2016 to 2018 at a center-of-mass energy of 13 TeV, corresponding to an integrated luminosity of 138 fb^{-1} . In addition, two projects for the high luminosity LHC upgrade are described: a highly parallelizable track-finding algorithm and managed networking for exabyte-scale data science.

Introduction

We are a way for the cosmos to know itself.

(Carl Sagan, *Cosmos: A Personal Voyage*, 1980)

Infinity in the palm of your hand

This dissertation is based entirely on a simple question: what is everything made of? For the vast majority of human history, all answers to this question were equally immaterial, yet for millennia, people with enough time on their hands and food in their bellies have wondered about it—from the Greek philosopher Democritus [1], to the Hindu scholar Kanāda (or, “atom-eater”) [2], and probably long before them. Over time, philosophical musing transformed into scientific inquiry, and as science progressed, faltered, corrected, and progressed again (and again, and again), patterns began to emerge. This evolution is presented here as what David Griffiths would call a “folk history,” with inspiration taken heavily from the first chapter of one of his revered textbooks [3].

In 1808, the chemist John Dalton published the first¹ atomic theory [4, 5], but the existence of the atom itself could not be directly verified. However, nearly a century later, physicists including Henri Becquerel, Marie Curie, and J.J. Thompson found that certain elements and electronic apparatuses emitted different kinds of invisible radiation [7, 8, 9, 10]. In the decades following, they further realized that this radiation consisted of tiny particles², and Ernest Rutherford directed rays of these particles at thin sheets of metal, probing the subatomic scale

¹Dalton’s now legendary claim to this was hotly contested by his contemporary, William Higgins, although without much credence [6].

for the first time. Based on these seminal experiments, Rutherford made a groundbreaking observation and two key deductions:

- *Observation:* The atom exists³ and has a very small “nucleus” of positive charge.
- *Deduction:* The nucleus is made of positively charged “protons.”
- *Deduction:* The nucleus is also made of “neutrons” which help keep it together.

Through their work in the lab, Rutherford, Curie, and the other atomic pioneers had unknowingly stumbled upon the surface of an even greater truth than Democritus, Kanāda, or anyone else could have imagined.

Albert Einstein, Max Planck, Niels Bohr, and others had begun writing down corrections to “classical” physics at the turn of the 20th century in order to better describe unexplained phenomena observed in the lab [14, 15, 16]. By the time Rutherford’s neutrons had been discovered by James Chadwick in 1932 [17], a cadre of theorists, including many of the most well-known physicists in history⁴, had formed around Quantum Mechanics. This “new physics” extended the early results of Einstein, Planck, and Bohr into a more complete theory describing the subatomic world, then thought to only include protons, neutrons, and electrons. Over the next decades, a slew of new particles were discovered: anti-matter was produced in the lab; electrons were joined by muons and neutrinos in the “lepton” family; an enormous crowd of more massive particles akin to protons and neutrons were observed (baryons and mesons which together comprise the families of hadrons). While some of these particles were already thought to exist, experimental results had, momentarily, outpaced theory. Theorists were developing new ideas on how to classify or relate different particles, but the ever-growing “particle zoo” was

²Namely photons (X-rays or γ rays), electrons (β rays), and helium ions (He^{2+} or α rays).

³It should be noted that Jean Perrin separately confirmed the existence of molecules, and thus reality of atoms, by proving experimentally that the random walk of a microscopic object suspended in liquid could only be due to collisions with the molecules of that liquid [11]—a phenomenon first observed by the great botanist Robert Brown [12] (i.e. Brownian motion) and rigorously explained by Albert Einstein in one of his first papers [13].

⁴Werner Heisenberg, Erwin Schrödinger, and Robert Oppenheimer to name a few.

otherwise in disarray. However, as is a consistent pattern in history, science progressed, faltered, corrected, and progressed again (and again, and again), and patterns began to emerge.

A glorious victory parade

It began with Murray Gell-Mann, who, among others, saw a pattern in the deluge of new particles being discovered. He arranged the numerous hadrons in geometric patterns—octagons, triangles, and so forth—according to their properties [18], and found that his unusual diagrams held predictive power: some of his shapes were missing vertices, which were later found to be particles exactly matching the expectations from his theory [19]. With confidence in Gell-Mann’s so-called “Eightfold Way” established, a natural question arose: what is the significance of these patterns? Gell-Mann himself, and, separately, George Zweig, arrived at the answer: hadrons are not themselves fundamental particles; instead, they are composed of truly fundamental particles called “quarks” [20], and the Eightfold Way diagrams corresponded to different configurations of quarks and the properties of those configurations. Similar to Rutherford’s experiments 50 years prior, wherein beams of particles were used to probe the structure of the atom, the reality of quarks could be probed by firing beams of electrons at protons, revealing their inner structure⁵ [21, 22]. And so the field progressed, with families of particles filling out into neat organizations: the quarks became a family of six, divided into three generations (up/down, charm/strange, top/bottom); the leptons, with the addition of the tau and its neutrino, satisfyingly also totaled six particles divided into three generations (e/ν_e , μ/ν_μ , τ/ν_τ); the photon was joined by the gluon, W, Z, and Higgs boson [23, 24], forming the family of bosons. At the same time, a beautiful, yet deeply perplexing, quantum field theory was developing, which describes the interactions between these particles. In the words of the great Sidney Coleman [25]:

This was a great time [1966 to 1979] to be a high-energy theorist, the period of

⁵These experiments really only showed that the proton was composed of “partons,” but they provided initial evidence for the existence of quarks.

the famous triumph of quantum field theory. And what a triumph it was, in the old sense of the word: a glorious victory parade, full of wonderful things brought back from far places to make the spectator gasp with awe and laugh with joy.

So began the Standard Model of particle physics: with groundbreaking discoveries and a scientific “victory parade,” which together still represent the largest fraction of Nobel Prizes by field [26].

There and back again

The Standard Model has, despite remaining largely unchanged since the 1980s, withstood a tremendous amount of scrutiny across decades of intense study and many billions of dollars of scientific instrumentation. It supplies, in principle, the answers to those questions posed at the beginning of this chapter: the origins and composition of everything in the universe are explained by the existence of, and interactions between, fundamental particles. However, we quickly found that this theory of everything is really only a theory of some things, and long-known errors and inconsistencies—reserved for discussion in the next chapter—have been, much to the community’s dismay, just as durable as the Standard Model has been otherwise. Particle physics has therefore reached a familiar inflection point, where we are left in anticipation of the next measurement that confirms an unusual approach or the next surprising discovery that leads the theory in an entirely new direction. Described in this document are the efforts of one graduate student, along with the invaluable support of his many colleagues, to contribute to this extraordinary enterprise. Together, we shall recount his journey, taking the perspective of the “experimentalist.” That is, very few formulas, axioms, and derivations—the language of theorists—are to be found in this dissertation. Instead, we will find tables of measurements, histograms, and some statistics. We will take general direction from the theory, then construct experiments to realize it in nature.

This work is based on the proton-proton collision data collected by the Compact Muon Solenoid (CMS), a multi-billion dollar particle detector situated on the beamline of the Large Hadron Collider (LHC)—the most powerful particle collider and the single largest scientific

instrument ever built, at the time of writing. At the risk of losing the reader who is not committed to reading the many pages that lie beyond this introduction, no new physics is documented here. Instead, this humble submission, amongst over one thousand publications from the CMS Collaboration alone, represents a small contribution to the rich history documented above, demonstrating the continued success of a theory that dares to describe the entire universe.

Chapter 1

A Beautiful Theory

What is especially striking and remarkable is that in fundamental physics a beautiful or elegant theory is more likely to be right than a theory that is inelegant.
(Murray Gell-Mann, Beauty and truth in physics, 2007)

1.1 The Standard Model

Put simply, the Standard Model (Fig. 1.1), shortened hereafter to “SM” for brevity, attempts to describe, at the most fundamental level, what the universe is made of and how it operates. So far, the content of the model is as follows: the universe is made of fermions (quarks and leptons), and it operates through the exchange of bosons (photons, gluons, W, Z, and H). That is, matter¹ is composed of assemblies of fermions, which are held together, pushed apart, and otherwise interact via the fundamental forces “transmitted” by bosons. The most familiar of these is the electromagnetic (EM) force, because its carrier, the photon, is absorbed by the retinas in our eyes, allowing sighted readers to review this text. Then, there is the weak force, carried by the W and Z bosons, which is responsible for the radioactive decay of certain elements. These first two forces are, in reality, understood to be unified into the “electroweak” (EW) force. Finally, there is the strong force, carried by the gluon, which holds the nuclei of atoms together. Alongside the forces, there is the Higgs mechanism, carried by the Higgs boson, which is responsible for endowing the fundamental particles with mass (Section 1.1.3). Together,

¹With the current exception of “dark” matter, which is discussed later in this chapter.

these forces and the Higgs mechanism—and gravity, whose omission here is left as a topic for another time—describe how *everything* came to be and continues to be: from the sight of the sun in the sky, to the nuclear fusion causing the sun to shine, to the formation of the sun and all of the other stars in the universe, to the first instance of creation itself. The entire universe is, we believe, the structure that emerges from the elegant, infinite dance of these fundamental particles.

Of course, a rigorous description of the SM could only be properly treated alongside a stack of textbooks. The concepts of fermions and bosons carry deeply insightful mathematics derived originally in statistical mechanics, which was itself developed to describe the behavior of large ensembles of microscopic objects (like gases). Moreover, the particle “zoo” has been herded into the confines of the mathematical framework called quantum field theory (QFT), wherein particles are described as excitations of quantum “fields” and physical laws are realized as mathematical symmetries. Unfortunately, four to five years of attentive instruction and rigorous study across multiple subjects—including quantum mechanics, special relativity, and a bevy of additional mathematical formalism—are required to even begin reading QFT textbooks, and well over a lifetime may be required to fully understand them². The SM will instead, as promised in the previous chapter, be described here from the perspective of an experimentalist: through hastily scrawled cartoons, rough calculations, and a great deal of hand-waving that, together, at least communicate the essential features of the theory.

1.1.1 Feynman diagrams

Fortunately, there is a way to encode essential SM calculations in simple drawings, so-called “Feynman diagrams,” which, consequently, help experimentalists keep track of physically allowed processes. In these pictures, time flows from left to right, while space is abstractly represented on the vertical axis³. The fundamental particles are represented by lines, and the intersections of three or four of these lines (Fig. 1.2) represent interactions between the

²Even still, paraphrasing the great Prof. Anesh Manohar, “all QFT textbooks are wrong.”

³In some dark corners of the physics community, these axes are switched, but this dissertation will not deviate from the configuration described here.

Standard Model of Elementary Particles

	three generations of matter (fermions)			interactions / force carriers (bosons)	
	I	II	III		
mass	$\approx 2.2 \text{ MeV}/c^2$	$\approx 1.28 \text{ GeV}/c^2$	$\approx 173.1 \text{ GeV}/c^2$	0	$\approx 125.11 \text{ GeV}/c^2$
charge	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	0	0
spin	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	0
	u up	c charm	t top	g gluon	H higgs
	d down	s strange	b bottom	γ photon	
	e electron	μ muon	τ tau	Z Z boson	
	ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	W W boson	

QUARKS (left side of the quark section)

LEPTONS (left side of the lepton section)

GAUGE BOSONS VECTOR BOSONS (bottom of the boson section)

SCALAR BOSONS (right side of the boson section)

Figure 1.1. The Standard Model of particle physics arranged in a table with the fermions placed in the left three columns (one per generation) and the bosons in the right two columns.

corresponding particles, so at least one of them must be a boson. For example, an electron emitting a photon (Fig. 1.3c) is represented by an electron coming in from the left, then turning into a photon and an electron leaving the picture to the right. This same vertex can be rotated clockwise, such that it instead depicts the annihilation of an electron and positron into a photon (Fig. 1.3a)—one of the electrons had to be replaced with its anti-particle (a positron) to conserve charge. Rotating it again, we see that it now represents an electron absorbing a photon (Fig. 1.3b). These vertices act as building blocks that can be rotated and fit together according to a set of rules that correspond to real physical laws. Thus, any fundamental physical process (e.g. Fig. 1.4)

can be represented with a Feynman diagram. Feynman diagrams are not only a visual aid for remembering which processes are allowed, however, for they also encode precise calculations about that process which can, importantly, be verified by experiment.




Figure 1.2. The fundamental vertices in the Standard Model. From left to right, top to bottom: (a) electromagnetic; (b), (c) electroweak; (d) weak; (e), (f), (g) strong; (h), (i), (j) Higgs. For compactness, a shorthand notation is used in some of the vertices to represent more than one particle: X^\pm is any charged particle, f is a fermion, and X_m is any massive particle.

1.1.2 Cross sections

Consider two electrons barrelling towards each other. In the classical picture, the electrons glance off one another and fly off to infinity at some angle to their original trajectories—like two errant ice skaters. This can be represented as a Feynman diagram (Fig. 1.5) which is assembled from two EM vertices. Again, time flows from left to right, showing the two electrons entering,




Figure 1.3. Rotations of the QED vertex, (a) an electron and positron annihilating and producing a photon, (b) an electron absorbing a photon, and showing (c) an electron emitting a photon.

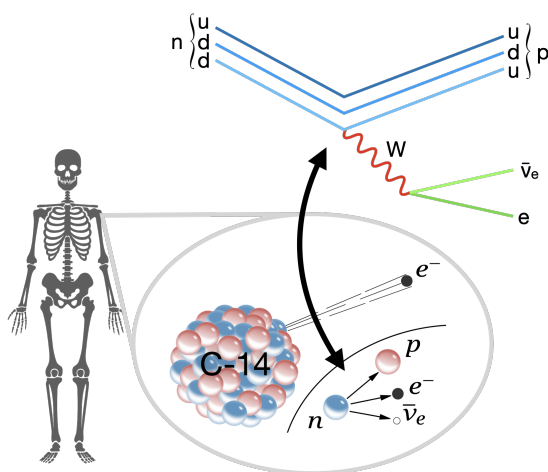


Figure 1.4. Feynman diagram for beta decay, which is the mechanism behind the radioactive decay of carbon-14 (C-14), which enables carbon dating.

then the exchange of a photon, followed by the two electrons leaving, much like the classical picture. Now, suppose we construct two beams of electrons, aim them at each other, then turn them on (preferably after we leave the room). In this scenario, we may well want to know the probability that the two electrons will bounce off each other (“scatter”). This probability is called the “cross-section,” because it is mathematically similar to the classical picture⁴, where we would compute the cross-sectional area presented to either of the colliding objects. Indeed, one of the most important features of QFT is the ability to compute the cross-section for scattering two electrons off of each other, as in this case, or any other interactions between particles. Feynman

⁴In fact, the answer to this question is expressed in the units of “barns,” literally as in “hitting the broadside of a barn,” coined by Manhattan Project scientists during World War II [27].

diagrams beautifully encode the complex mathematics at work here: each line and every vertex (where lines meet) correspond to a term in the calculation. For the electron-electron scattering example, this would look like

$$\langle |\mathcal{M}|^2 \rangle = \frac{2g_e^4}{(p_1 \cdot p_3)^2 (p_1 \cdot p_4)^2 [(p_1 \cdot p_2)^4 + (p_1 \cdot p_3)^4 + (p_1 \cdot p_4)^4]} \quad (1.1)$$

where p_1 and p_2 are the four-momenta of the incoming electrons, p_3, p_4 are the four-momenta of the outgoing electrons, and g_e is the EM coupling constant. By integrating the “amplitude” \mathcal{M} over the phase space—effectively, the different ways one can apportion the available energy amongst the particles in the final state of the system—one can finally obtain the cross-section. This exercise, and the compact answer above, serves to demonstrate the sheer complexity of QFT calculations. Many details were left out beyond the computation itself, including details surrounding the spin of the electron (e.g. the quoted amplitude is “spin-averaged”), the definition of spin, the definition of four-momentum, all of special relativity, and so on. Those details, and much more, can be found in Chapter 7 of the illustrious textbook from which this entire example was borrowed [3] alongside a few years of physics education at an accredited educational institution. Moreover, this feature of QFT—the ability to compute cross-sections—is absolutely vital because it offers a way of testing the theory with observation: compute the probability that the event is expected to occur, then try to reproduce that event many times in the lab and see how many times it really happens.

The validity of any model is a precarious condition: the model must exactly describe reality, else it is not a realization of the truth, but rather only approximately—or worse, accidentally—correct. That is, *every* SM cross-section value must be accurate for the validity of the model to hold. And so generations of physicists have made diligent, and highly accurate, measurements of dozens of cross-sections—spanning orders of magnitude of rarity. Tabulated in grand tables (Fig. 1.6 and 1.7), it is evident that the SM has so far withstood every test derived from generations of experiments. This is an incredible feat: a “simple” model seems to describe




Figure 1.5. Leading-order Feynman diagram for electron-electron scattering.

subatomic physics, which drove the creation of the universe and continues to drive the mechanics of *everything* in existence, with incredible accuracy. Lingering beyond these shining trophies, however, are a number of glaring inconsistencies and enormous missing pieces. In other words, we know with certainty that the SM is not complete, and that much more work is needed to complete it. We will focus in this dissertation on questions around the Higgs boson. However, we will also mention some additional open questions to illustrate the magnitude of the problem ahead for future particle physicists.

1.1.3 The Higgs boson

Because we swore to steer away from the details of QFT, we cannot define the Higgs boson and Higgs mechanism completely. For this, we are better served by a primer by Griffiths [3], followed by the standard QFT textbooks [30, 31], and finally the original papers on the subject [32, 33, 34, 35]. Without these prerequisites, it suffices to say that the idea of the Higgs boson is deeply based in Lagrangian mechanics. Originally derived for classical systems—an arrow in flight, a planet in motion, a ball on a ramp—an appropriately written “Lagrangian,” a mathematical object, together with the Euler-Lagrange equation can reproduce the mathematical description of the dynamics of a given system—the position of the arrow, planets, or ball as a

function of time. The Lagrangian L is a simple function of the energies in a classical mechanics:

$$L = \underbrace{\frac{1}{2}m\dot{x}^2}_{\text{kinetic energy}} - \underbrace{U(x)}_{\text{potential energy}} \quad (1.2)$$

where x , really $x(t)$, is the position of some physical object in one direction as a function of time and \dot{x} is shorthand for the derivative of x with respect to time ($\frac{d}{dt}$), or the velocity. By plugging L into the Euler-Lagrange equations,

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{x}} \right) = \frac{\partial L}{\partial x} \quad (1.3)$$

we can solve for $x(t)$ with some algebra and calculus, trivializing some of the most challenging undergraduate classical mechanics problems.

In particle physics, we promote $x(t)$ to a field⁵ ϕ , which is a function of space and time $\phi(x, y, z, t)$, and L to \mathcal{L} , a Lagrangian *density*. Our interest in these abstract fields is well-motivated in QFT, where a field $\phi(x, y, z, t)$ with certain properties corresponds to an observable particle with those same properties. For example, the Lagrangian

$$\mathcal{L} = \underbrace{\frac{1}{2}(\partial_\mu \phi)(\partial^\mu \phi)}_{\text{kinetic term}} - \underbrace{\frac{1}{2}\left(\frac{mc}{\hbar}\right)^2 \phi^2}_{\text{mass term}} \quad (1.4)$$

represents a spin-0 particle with mass m . As in classical mechanics, plugging the appropriate Lagrangian density into the Euler-Lagrange equation yields a precise description of the dynamics of fundamental particles.

The Higgs mechanism arises when we try to write down a Lagrangian density \mathcal{L} that accounts for the existence of the W and Z bosons. In doing so, we find that we must demand that

⁵Fields are used here in a quantum context (scalar fields in particular) but there are many classical examples: the temperature of a room, the height of the sea, the strength of a magnetic field, i.e. quantities that may be different at every position in space and may also vary in time.

\mathcal{L} does not change under certain transformations of the fields ϕ , which, at face value, equates to demanding that the field is massless—the finer details here are far beyond the scope of this dissertation. This is a problem because the W and Z bosons are verifiably not massless. However, by introducing two scalar fields—corresponding to the Higgs boson and Goldstone boson—and massaging the mathematics, a mass term emerges for the W and Z bosons. The Goldstone boson does not correspond to a real particle, but instead services the mathematics⁶ and disappears under a specific transformation. The Higgs *mechanism* itself refers to the appearance of the W and Z boson masses—also the fermion masses, but through different interactions [36, 37].

Beyond the blackboard, the Higgs boson gives us a picture of how the universe came to be: in the first moments of time, the universe cooled, allowing the Higgs field to saturate all of space, and the fundamental particles were thereby endowed with mass by the Higgs mechanism. Thus, rather than fly off to infinity at the speed of light, they formed atoms, and the universe as we know it bloomed. The Higgs boson is therefore essential to the origin and continued existence of the known universe—in fact, if the Higgs field were to spontaneously disappear, the entire universe would evaporate in a nanosecond [38]. It is also, however, a key to understanding the universe’s distant future [39]: will the universe collapse or explode or something else entirely? We will elect to spare ourselves of existential crisis and, instead, simply appreciate the specific importance of the Higgs boson. While every piece of the SM is important, like links in a chain, the Higgs boson is also one of the less well-understood particles, so further study of this latest addition to the SM could shed crucial light on the many open questions in particle physics.

1.2 Selected open questions

1.2.1 What is dark matter?

In the 1970s, Vera Rubin, among others, was making careful observations of the rate at which distant galaxies spin on their axes [40]. By also recording the light emitted by each

⁶For example, it accounts for the fact that there are three bosons that carry the weak force (W^+ , W^- , Z).

galaxy, she could estimate the distribution of mass within it, and therefore calculate the expected rotational velocity due to the gravitational pull of a galaxy's constituents. The observation and prediction did not align, leading us to believe that there is some “dark” matter that we cannot see and is therefore missing from our calculations. In fact, based on the observable mass in a galaxy, we would expect it to fly apart [41]. Today, as we probe the subatomic world at unprecedented scales, we have yet to observe any new particle that might be the right candidate for dark matter—really, we have yet to observe any new particle whatsoever. Dark matter is therefore an enormous missing piece—27% of the universe [42]—in the SM.

1.2.2 What is dark energy?

Cosmology yields yet another missing piece in the SM, but at a completely different scale than dark matter. Measurements of the cosmic microwave background (CMB) radiation, light from the early universe, and of distant galaxies shows clear indication that the universe is expanding: all of this light has been stretched out, with visible light appearing an eerie red. But this expansion is accelerating, and we do not know why. More precisely, we are left wondering: where does the enormous amount of energy required to drive this acceleration come from? With no known answer to this question, we have yet another invisible quantity—this time, a “dark” energy that constitutes 68% of the universe [42]—as a placeholder.

1.2.3 What is the true nature of the Higgs boson?

The Higgs boson holds many secrets yet to be uncovered, despite being discovered over a decade before the publication of this dissertation. It is one of the central physics targets of the LHC for the coming decades, and only a few of the most interesting questions are summarized here:

- *Is the Higgs boson really a fundamental particle?* Pursuing an ever-reductive line of reasoning, it is important to always ask if a particle is in fact a composite particle, rather than a fundamental one. A number of theories that predict a composite Higgs boson exist,

e.g. Refs. [43, 44, 45], and, if true, they may explain other open questions around the nature of the Higgs boson.

- *Is there only one Higgs boson?* In a 2012 interview at CERN, only 6 months before the discovery of the Higgs boson, Murray Gell-Mann mused

As to there being one boson that results from this Higgs process, it's not quite so certain because there could be more than one... and [CMS and ATLAS] might not be looking for it in the right way for that.

Although only a single Higgs boson was discovered, the existence of others, e.g. Refs. [46, 47, 48], could again answer other open questions around the nature of the Higgs boson.

- *Why is the Higgs boson mass so small?* While nearly as heavy as the most massive SM particle, the top quark, the Higgs boson is not nearly as massive as it “should” be. Corrections from heavier particles—even unknown ones—could make the Higgs boson over 16 orders of magnitude heavier, so some new physics should correct this somehow [49, 50].

Answers to many of the questions above could lie in more precise measurements of the properties of the Higgs boson, such as those presented in this dissertation.

Overview of CMS cross section results




Figure 1.6. The totality of the cross-section measurements performed with data from the CMS Experiment compared to the Standard Model predictions, from Ref. [28]. The size of each cross-section is plotted on the x-axis and each measurement is separated along the y-axis, with the process name, collision energy, and journal reference on the left and the integrated luminosity of the collision data used in the measurement on the right. Precise agreement with the Standard Model can be seen across several orders of magnitude, representing the triumph of the model across decades of experimental scrutiny.

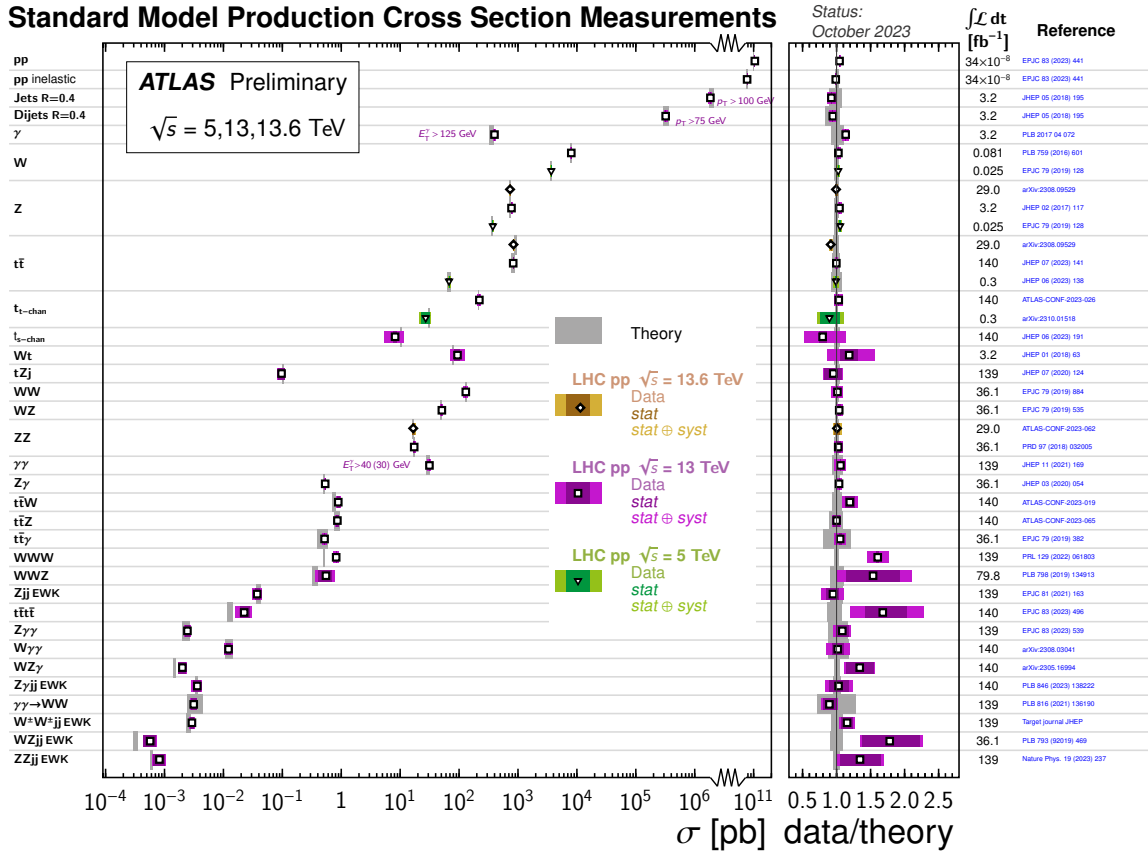


Figure 1.7. A selection of cross-section measurements performed with data from the ATLAS Experiment compared to the Standard Model predictions, from Ref. [29]. The size of each cross-section is plotted on the x-axis and each measurement is separated along the y-axis, with the process name on the left and the integrated luminosity of the collision data used in the measurement and the journal reference on the right. Precise agreement with the Standard Model can be seen across several orders of magnitude, representing the triumph of the model across decades of experimental scrutiny.

Chapter 2

A Grand Apparatus

A glance at the ATLAS and CMS detectors at CERN reveals their beauty... These detectors are the modern cathedrals of the rational world created by scientists, experimentalists, and theoreticians.

(François Englert, CMS: The Art of Science, 2016)

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is the largest particle accelerator ever built. Its most striking feature is a 27 km ring buried 100 m beneath the Franco-Swiss border (Fig. 2.1) in which two beams of protons (or heavy ions like lead), after going through several smaller stages (Fig. 2.2), are accelerated to 99.9999991% of the speed of light in opposite directions. The beams are steered by thousands of magnets, including 1232 superconducting dipole magnets (Fig. 2.3), placed along the circumference of the ring. At various points, the proton beams are directed towards each other, allowing the protons to collide. These collision points are surrounded by enormous, multi-layered particle detectors which record snapshots of the collisions.

The protons are accelerated in bunches, composed of approximately 115 billion protons each, so when the bunches are brought together (called a “bunch crossing”), over 200 billion protons are brought very close together. However, only a small portion of them actually interact. To increase our odds of producing something truly interesting, the bunch crossings are spaced close together, with only 25 nanoseconds of separation. To put this into context, the speed of




Figure 2.1. A diagram illustrating the depth of the LHC beneath the surface (left), from Ref. [51], and an aerial photograph of the entire CERN complex (right), from Ref. [52]. The SPS and LHC tunnels illustrated in both figures.

light is roughly 1 ft/ns, so the next bunch crossing will occur before the light from a screen on one side of the CMS control room can reach the eyes of someone standing on the other side (approximately 39 feet [55]).

2.1.1 Collision energy

The collision energy, often written as \sqrt{s} , is measured in electronvolts (eV), the standard unit of energy in particle physics—at LHC scales, teraelectronvolts (TeV), or trillions of eV, are the relevant units. In Run 1 of the LHC (2009–2013) the collision energy was ramped up to a maximum of 8 TeV, much lower than what the LHC was designed for. This was done out of an abundance of caution due to the fact that the LHC exploded¹ when it was first turned on in 2008.

Then, in Run 2 (2016–2018), the collision energy was increased to 13 TeV. Over the course of Run 2I, the collision energy will be ramped up to 14 TeV, where it will be held for the remainder of the LHC’s lifetime. At 7 TeV per proton beam, each bunch has roughly 7 times the kinetic energy of a flying mosquito with a mass ten trillion times smaller.

¹A faulty cable lead to a catastrophic quench of one of the LHC superconducting magnets, resulting in the release of several metric tons of helium (used to cool the magnets) into the LHC tunnel [56].

The CERN accelerator complex *Complexe des accélérateurs du CERN*

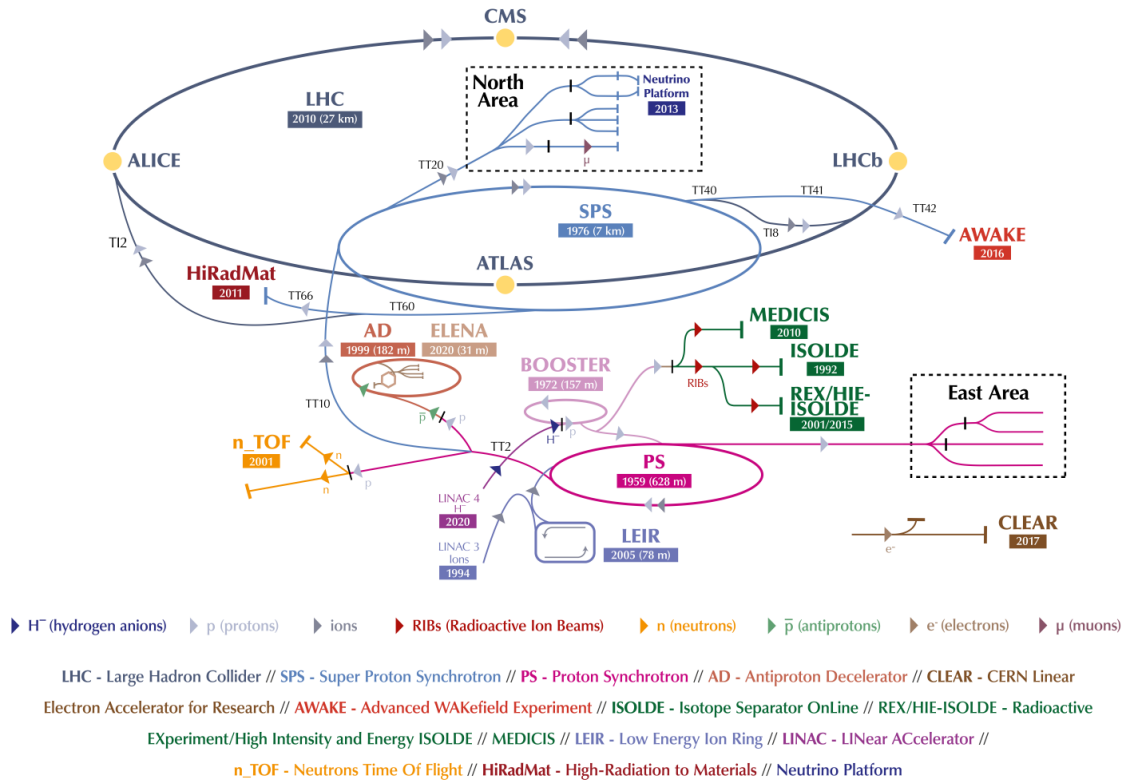


Figure 2.2. The CERN accelerator complex, from Ref. [53], illustrated in detail. The different stages of particle acceleration can be seen in detail, described here for protons. First, negative hydrogen ions (H^-) generated by LINAC 4 are fed into BOOSTER, which strips the electrons from the H^- ions, leaving only the protons, and accelerates them to 2 GeV. Next, the Proton Synchrotron (PS) accelerates the protons to 26 GeV. The PS feeds into the Super Proton Synchrotron (SPS) which further accelerates them to 450 GeV. Finally, the protons are fed into the LHC, which accelerates them to a final energy of 7 TeV.

2.1.2 Luminosity and pileup

The luminosity is a measure of an accelerator’s ability to initiate the desired particle interactions [57]. It is thus measured in the inverse units of a cross-section, namely cm^{-2} or inverse barns (b^{-1})—or, at LHC scales, inverse femtobarns (fb^{-1}). At the LHC, luminosity is a measure of the ability to initiate pp collisions. Immediately, it is clear that it should be some function of time: the longer we let the beams collide, the more pp collisions produced.




Figure 2.3. A cutaway diagram illustrating the internals of a dipole magnet inside the LHC tunnel (left), from Ref. [54], and a photograph of a decommissioned dipole magnet with a physicist for scale (right).

This is called the *instantaneous* luminosity (\mathcal{L}), or how many interactions are initiated at any one time. We can also consider the *integrated* luminosity, or the total number of interactions initiated over some period of time—as the name suggests, this is simply the integral of the instantaneous luminosity over time. Put simply, the integrated luminosity represents the rate at which interactions occur at the LHC, while the integrated luminosity is a compact shorthand for how many interactions occurred at the LHC during some period of data-taking. For example, with the LHC operating at an instantaneous luminosity of $\mathcal{L} = 1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, given the pp inelastic scattering cross-section at 13 TeV ($\sigma = 71.3$ millibarns [58]), we can expect $\mathcal{L} \times \sigma = 713$ million pp collisions per second, or 11.24 quadrillion pp collisions over 6 months. This is notably much larger than 1 pp collision per bunch crossing: with one bunch crossing every 25 ns, one might expect a mere 630 trillion collisions. In reality, under Run 2 conditions, there are approximately 30 pp collisions per bunch crossing on average (Fig. 2.4), explaining the difference between these back-of-the-envelope calculations, and out of all of those collisions, only one (if any) pp collision will initiate a process that a physicist at the LHC wants to study (e.g. Fig. 4.1). The rest, called “pileup” collisions, each produce some number of low-energy particles, presenting a sizable amount of noise for extracting the interesting physics. Importantly, a cross-section σ for a specific process—the cross-section for a Higgs boson to be produced in a

pp collision, for instance—can be used in the same calculation to determine how much of it is produced at the LHC.

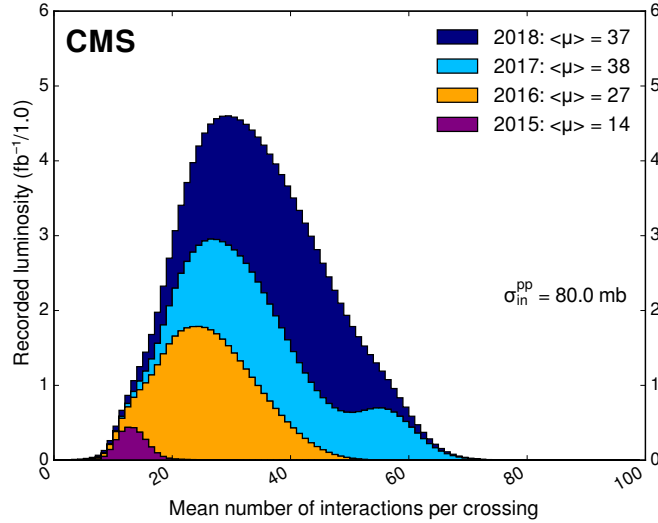


Figure 2.4. Distribution of the average number of interactions per bunch crossing within the CMS detector from Ref. [59]. The distributions for 2015, 2016, 2017, and 2018 are stacked together, showing the composition of the overall distribution.

Based on intuition, we can guess that the luminosity must be a function of the number of protons in a bunch, the bunch spacing, and the exact dynamics of the bunch crossing. Indeed, at the LHC, we define the instantaneous luminosity as follows:

$$\mathcal{L} = \gamma \frac{n_b N^2 f_{\text{rev}} R}{4\pi \beta^* \epsilon_n} \quad (2.1)$$

where n_b is the number of bunches per beam, N is the number of protons per bunch, and f_{rev} is the frequency at which bunches revolve around the LHC ring, while γ comes from special relativity and β^* , R , and ϵ_n represent physical properties of the proton beams [60]. Therefore, in order to modify the rate at which interactions may occur, we have several parameters that we can tune.

2.2 The Compact Muon Solenoid

The Compact Muon Solenoid (CMS) Experiment is one of two general purpose LHC experiments, the other being the ATLAS² Experiment, among the four major experiments supported by the LHC [61], where the other two are more specialized: ALICE, for studying heavy ion collisions, and LHCb, for studying b quarks. Compared to ATLAS, which stands at a mighty $46 \times 25 \times 25$ meters in dimension, CMS is “compact” at a stout $21 \times 15 \times 15$ m, with a dedicated muon system and one of the world’s largest solenoids [62, 63]. See Fig. 2.5 for its exact specifications, and Figures 2.6 and 2.7 for some beauty shots.




Figure 2.5. A detailed cutaway diagram of CMS, from Ref. [64], with each subdetector labeled with its name and some of its characteristics.

²Whereas ATLAS originally derived its name from, putting it politely, a rather creative acronym: “A Toroidal LHC Apparatus.”




Figure 2.6. CMS in all of its glory, with one of the endcaps separated from the main body of the experiment, pictured from the top (left), from Ref. [65], and side (right), from Ref. [66].




Figure 2.7. CMS in the closed configuration, pictured from the top (left), from Ref. [67], and side, with a physicist in the foreground for scale (right).

2.2.1 Overview

The CMS Experiment is composed of subdetectors arranged in consecutive layers surrounding one of the LHC collision points, where each layer interacts with or completely absorbs certain particles, producing an electric signal that can be used to measure some property of those particles. The innermost layer is the silicon tracker, which allows for the reconstruction of the trajectories of throughgoing charged particles (“tracks”). Next is the electromagnetic calorimeter (ECAL), which absorbs electrons and photons and records their individual energies. After the ECAL, there is the hadronic calorimeter (HCAL), which absorbs hadrons and records their individual energies. These first three layers—the tracker, ECAL, and HCAL—are surrounded

by the eponymous CMS magnet, a superconducting solenoid that immerses the inner layers in an approximately uniform magnetic field that runs parallel to the beamline. This is critical: charged particles fly along curved trajectories in a magnetic field according to their charge and momentum, so those properties can be inferred from a high-quality measurement of each particle's trajectory. Finally, there are alternating layers of muon chambers, the other half of the experiment's namesake, and iron support structures. The former detects throughgoing muons, which pass through all of the inner layers mostly unperturbed, and measures an additional portion of their tracks. However, the latter is equally important: the iron "return yoke" guides the magnetic field outside of the solenoid, absorbs stray particles that make it past the inner layers, and supports the immense weight of CMS itself. By combining information from all of these detectors, the exact identity of any individual particle can, in principle, be inferred based on which detectors registered a signal. Therefore, a full "picture" of each pp collision event is recorded by CMS for further study. The exact function of each subdetector layer described here is detailed below.

2.2.2 Superconducting solenoid

The curve of a track is critical, as it allows us to infer the charge and momentum of a particle. However, in a weak magnetic field, the particles produced at the LHC would have nearly straight tracks, due to the large pp collision energy. The magnetic field inside CMS must therefore be very large [68]. It should also be nearly uniform everywhere, in order to make the determination of each particle's charge and momentum as simple as possible. The CMS magnet must also be large in dimension, however, as it must surround the tracker, ECAL, and HCAL, since it would otherwise block outgoing particles.

By winding copper wire into a helix and passing a current through it, we can generate a magnetic field whose strength is directly proportional to the current and number of turns of the wire, but inversely proportional to the length of the helix. Within the volume of the helix, the magnetic field will be almost uniformly oriented in a single direction, determined by the

orientation of the helix and direction of the current (Fig. 2.8a). This is not true outside of the helix, where the magnetic field lines curve in space.

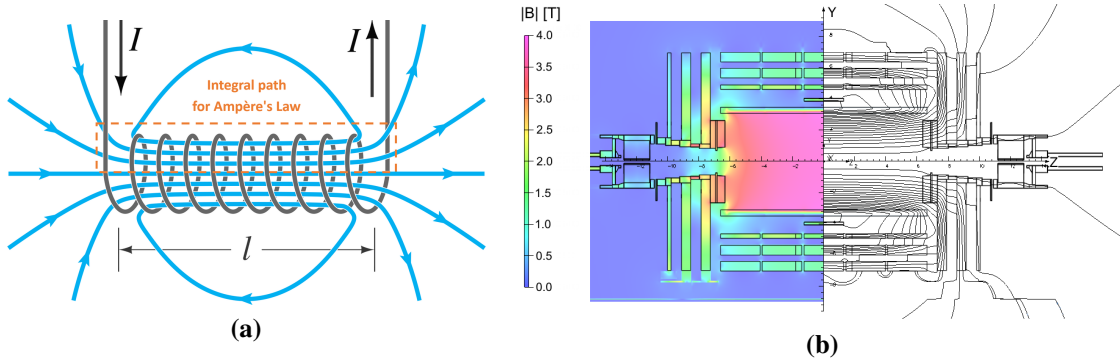


Figure 2.8. The field lines of an ideal solenoid (a), from Ref. [69], and those of the CMS magnet (b), from Ref. [70]. For the ideal solenoid, the magnetic field is described by a simple equation: $B = \mu_0 \frac{NI}{l}$, where μ_0 is the magnetic constant, N is the number of turns, I is the current, and l is the length of the helix.

The CMS magnet [71] is a massive³ realization of a solenoid. It is composed of over 2000 turns of Rutherford wire, which has a rectangular cross-section (Fig. 2.9). In operation, it is cooled to superconducting temperatures (-268.5 °C, or one degree warmer than outer space), such that a high current (20 kiloamperes⁴) can be passed through the coil without violently destroying it. This configuration allows the CMS magnet to produce a mostly uniform magnetic field of 4 T inside the solenoid volume. That is 2–4 times larger than an MRI, which are typically 0.5–1.5 T [73], within a volume that is 1000 times larger. We must also, however, have a fairly uniform magnetic field outside of the solenoid in order to maintain good momentum resolution for muons. To achieve this, the iron support structures that hold CMS together are also designed to guide the magnetic field lines (Fig. 2.8b). Inside the iron, the magnetic field runs almost uniformly parallel to the field inside the solenoid, but in the opposite direction—this gives muon tracks a characteristic S-shape in the transverse plane (Fig. 2.15b).

³It was built offsite, however, so it needed to be designed to fit within 7 such that it could be wheeled through the streets in Cessy, France (under which CMS is situated) [68].

⁴For scale, this is 100 000 times larger than the minimum fatal current for humans (100–200 milliamperes) [72].




Figure 2.9. A photograph of a section of the CMS magnet, from Ref. [74], showing the cross-section of the solenoid winding (left), and a diagram of the cross-section of a Rutherford cable, from Ref. [71] (right).

2.2.3 Silicon tracker

In order to reconstruct the track of any particle, we need its position at multiple points in time. A tracker records these positions. Mechanically, this can be done in a variety of ways. Some of the earliest particle physics experiments used bubble chambers, a device that maintains (for a brief instance) a volume of superheated liquid immersed in a magnetic field in which throughgoing charged particles leave helical trails of bubbles. Photographs of these trails were used to infer the identities of each particle. Still others used even more exotic non-electric solutions, like the OPERA Experiment, which used enormous layers of nuclear emulsion film (modified photography film) in which throughgoing particles would leave tiny black dots after the film was developed⁵ [75]—the experiment used over 100 000 m² of emulsion film in total.

The leading challenge in designing the CMS tracker was the unprecedented level of radiation, due to the high frequency of pp collisions, which each produce many thousands of

⁵Developing the film was quite an operation, and a large robotic arm was used to remove and replace the layers.

particles. Existing tracker designs, like the CDF Experiment’s beautiful gas-and-wire tracker [76], would quickly break down⁶ in this kind of environment. Thankfully, advances in radiation-hard electronics yielded a new solution. In particular, a silicon-based tracking module was devised. When charged particles pass through the module, they liberate electrons from the silicon atoms. These electrons are collected on one of many small readout plates, which results in a localized electric signal called a “hit,” giving excellent spatial resolution.




Figure 2.10. A photograph of the CMS silicon tracker inner forward pixel modules (left), from Ref. [78], and the outer barrel layers with a physicist in the background for scale (right), from Ref. [79].

The CMS tracker is composed of multiple layers of silicon tracker modules (Fig. 2.11), such that each throughgoing particle leaves multiple hits that can then be used to reconstruct its track. The layers are grouped into different sections according to their proximity to the beamline and their geometry. The innermost layers comprise the “inner tracker” (Fig. 2.10a). In these layers, pixel modules are used because they have superior spatial resolution, allowing us to measure the first portion of each track precisely. The outermost layers comprise the “outer tracker” (Fig. 2.10b). These layers use strip modules, which have worse spatial resolution than the pixel modules, but are easier to produce.

⁶In fact, the tracker section was left blank in the first CMS design document, since there were no known solutions at the time [77].




Figure 2.11. A detailed cutaway diagram of the CMS Phase-1 tracker from Ref. [80].

2.2.4 Electromagnetic calorimeter (ECAL)

The ECAL (Fig. 2.12a) consists of nearly 80 000 lead tungstate (PbWO_4) crystals⁷ (Fig. 2.12b) [81]. When an electron or photon passes through one of these crystals, they interact through the electromagnetic force with the atoms in the crystal lattice, resulting in the emission of photons. This process is called “scintillation” and the amount of light produced from one of these interactions is proportional to the energy of the impinging particle. Each PbWO_4 crystal is glued to a photodetector that collects the scintillation light by the attached crystal and outputs an electric signal proportional to the amount of light collected. This signal can then be used to measure the energy of electrons and photons.

The probability of material interactions occurring as well as the shape of the shower varies amongst scintillating materials. Ultimately, PbWO_4 was selected for its short radiation

⁷These crystals were grown in manufacturing plants in Russia and China and took roughly a decade to produce [81].

length and small Molière radius [82], corresponding to a high probability of interactions and showers localized mostly to a single crystal, respectively. The former property is most obviously useful: a high probability of scintillation means electrons and photons are more efficiently collected by the ECAL. The latter property, however, is also vital: localized showers are less likely to overlap, so individual particles can be resolved.




Figure 2.12. A photograph of the CMS electromagnetic calorimeter (ECAL) after half of the modules had been installed (left), and a physicist posed with the lead tungstate scintillator crystals used in the ECAL (right). Both photographs are from Ref. [83].

2.2.5 Hadronic calorimeter (HCAL)

The HCAL (Fig. 2.13a) is composed of interwoven layers of solid brass or steel plates and plastic scintillator tiles [84]. When a hadron interacts with one of the brass and steel “absorber” layers, a shower of secondary particles are created. The secondary particles pass through the next scintillator layer, creating scintillation light, then proceed to the next absorber layer. In this way, large cascading showers of secondary particles are produced according to the energy of the impinging particle, where each scintillator layer gives some measurement, or “sample,” of the original particle’s energy.

This kind of calorimeter, which has alternating absorber and scintillator layers, is called a “sampling” calorimeter as opposed to a “homogeneous” calorimeter, wherein a single medium is used to both create the shower and measure the energy. Sampling calorimeters are typically less

costly and give some information on the shape of the shower, whereas homogeneous calorimeters have superior energy resolution. However, a homogeneous hadronic calorimeter would need to be infeasibly large in order to fully contain a hadronic shower [85]. It is particularly critical that the HCAL be compact for CMS, as every increase in size of the inner layers results in large increases in volume of the solenoid and muon chambers, resulting in enormous increases in energy and material costs. Due to these space constraints, an additional layer of the HCAL needed to be placed outside of the magnet in order to measure, and absorb, anything that made it past the bulk of the HCAL. In fact, the HCAL was already projected to be too costly due to the sheer amount of high-quality brass needed to create the absorber plates. To save on costs, the Russian Navy was convinced to recycle over one million World War II artillery shells (Fig. 2.13b), which were designed to withstand a large amount of stress (read: explosions) and years at sea, for the cause [86].



Figure 2.13. A photograph of the CMS hadronic calorimeter (HCAL) in front of the magnet and muon chambers before the detector was installed in the experiment cavern (left), from Ref. [87], and workers in Mormansk sitting on the decommissioned World War II artillery shell casings used to supply the brass for the HCAL (right), from Ref. [88].

2.2.6 Muon chambers

The CMS muon system (Fig. 2.14a) was originally composed of three different kinds of detectors: drift tubes (DTs) in the barrel, cathode strip chambers (CSCs) in the endcaps, and resistive plate chambers (RPCs) interwoven between the layers of DTs or CSCs. All of the

different types of muon chambers operate on a similar principle: muons pass through a chamber filled with gas, liberating electrons from the gas atoms; those electrons are then collected on a conductive surface, resulting in an electric signal [89]. The DTs (Fig. 2.15) and CSCs are assembled into multiple layers that serve as an enormous outer tracker (Fig. 2.14b) exclusively for muons. DT and CSC modules also both have multiple inner layers in order to give a measurement of a muon’s outer track accurate enough to be matched to its inner track in the silicon tracker. The RPCs, meanwhile, are designed to give a much faster, if less accurate, measurement of each muon’s momentum. Because the presence of muons indicates potentially interesting physics (most collisions at the LHC simply produces low-energy quarks), a signal in the RPC can be used to determine whether a pp collision should be recorded. Moreover, the time resolution (2 ns) of the RPCs makes this information available well before the next pp collision.

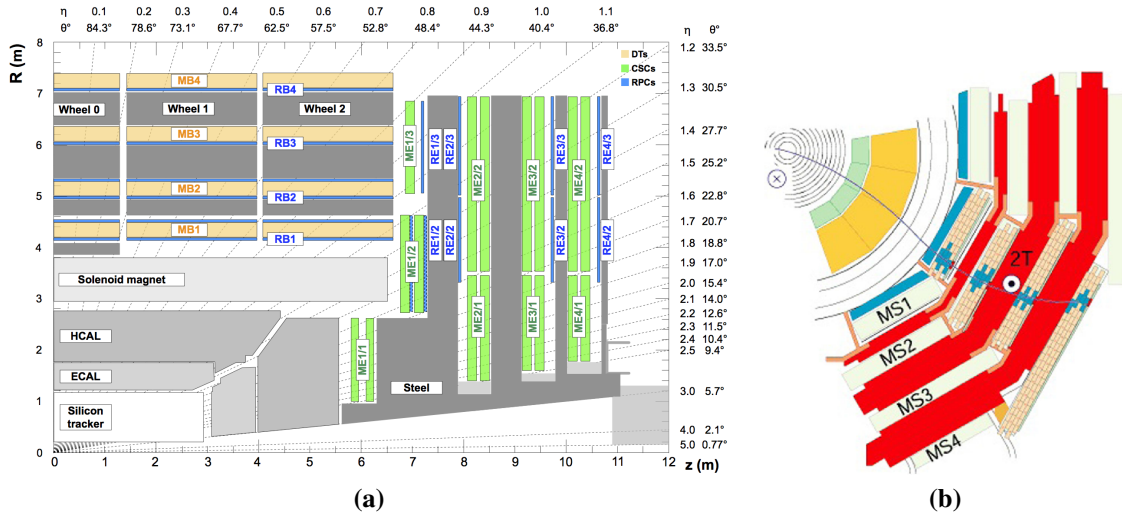


Figure 2.14. The cross-section of CMS in the r - z plane (left), from Ref. [90], and r - ϕ plane (right), from Ref. [91]. In the r - z plane, the DTs, CSCs, and RPCs are drawn in light green, dark green, and blue, respectively, showing the layout of the muon system as well as the other subdetectors. In the r - ϕ plane, the muon “stations” in the barrel are depicted with the DT internals exposed for stations with a throughgoing muon, showing the activation of the individual DT cells (blue) as well as the activation of a section of the RPC (red).

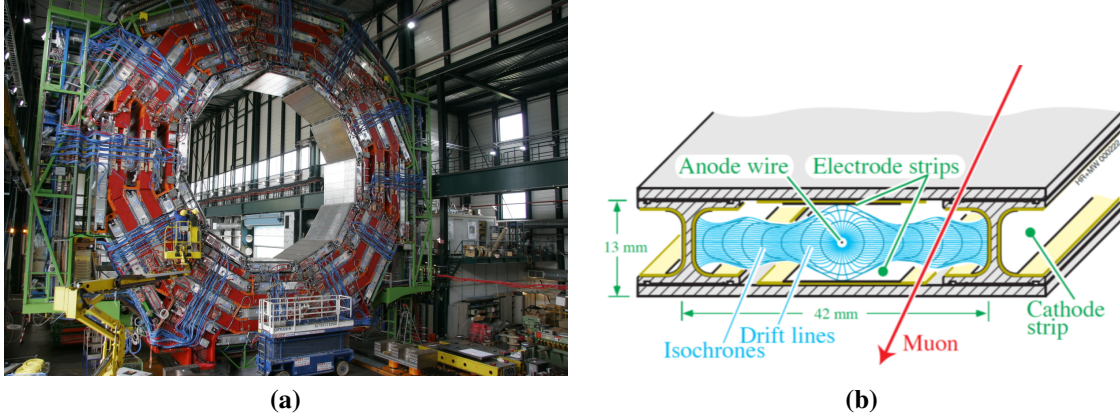


Figure 2.15. A photograph of a section of the barrel muon chambers (drift tubes) before they were installed (left), from Ref. [92], and a diagram of a drift tube muon detector (right), from Ref. [93].

2.2.7 Trigger system

If every bunch crossing was recorded, with one bunch crossing every 25 ns (40 MHz), the CMS Experiment would record data at an intractable rate:

$$40 \text{ MHz} \times 0.75 \text{ megabytes per event} = 30 \text{ terabytes per second} \quad (2.2)$$

Within 24 hours, CMS alone would produce over 2 *exabytes* of data—within a week, CMS could thus easily overwhelm CERN’s storage capacity. However, pp collisions rarely produce something worth studying like a Higgs boson, vector boson (W or Z), or something new, and anything new had better be rare, else we should have seen it already. We can therefore discard most pp collisions, making the continued operation of CMS sustainable.

The CMS Experiment has a two-stage trigger system for identifying interesting events. The first stage, called the “Level 1” (L1) trigger, takes raw data from the calorimeters and muon chambers and makes a decision on whether to keep or discard the event. It is implemented with dedicated hardware, so it can run synchronously with the pp collisions. Although it can only make very simple decisions, the L1 trigger reduces the output of CMS from 40 million events per second down to 100 000 by discarding events that are obviously uninteresting. The second

stage, called the “high level” trigger (HLT), relies on high-throughput software (written in C++), rather than specialized hardware, to reduce the event rate from 100 000 events per second to 1000 events per second. The HLT is designed to ingest more sophisticated input and thus make more precise decisions. In particular, events are processed along trigger “paths,” representing different sets of selections. Passing events are then sorted and written to disk for further processing.

2.3 The high luminosity era

There are three primary knobs to turn at the LHC: what things are being collided, what energy they are collided at, and how many of them are collided at once (the instantaneous luminosity). Since most of the experiments serviced by the LHC were designed for pp collisions, and because the LHC is designed for at most $\sqrt{s} = 14\text{ TeV}$, the final era of LHC physics will see a massive increase in luminosity. That is, at the “high luminosity” LHC (HL-LHC), there will be 100–200 concurrent pp collisions per bunch crossing (Fig. 2.16b), corresponding to an increase by a factor of 5–7.5. More collisions mean more opportunities for interesting physics. For example, in 2017, the LHC produced 3 million Higgs bosons per year, whereas the HL-LHC will produce over 15 million per year [94].

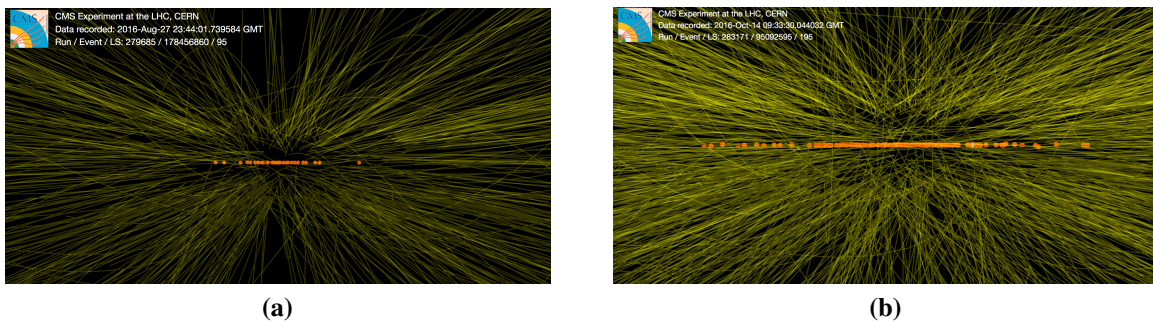


Figure 2.16. A collision event with 30 reconstructed vertices amidst high pileup conditions for Run 2 recorded at CMS in 2016 (a) is shown next to an event with 86 reconstructed vertices amidst HL-LHC-like pileup recorded at CMS in the same year (b), from Refs. [95, 96]. The dots are reconstructed vertices and the thin lines are the reconstructed particle tracks.

Significant modifications to the LHC as well as its experiments are needed to enable

HL-LHC operations. The CMS Experiment, in particular, will receive several upgrades:

- *Enhanced silicon tracker* [97]: The tracking modules in the outer tracker will be replaced with bilayer modules that give an approximation of a throughgoing particle's p_T . These modules will give the L1 trigger new information, as well as provide an opportunity for an entirely new paradigm for offline track reconstruction (Ch. 6).
- *Timing detector* [98]: A new detector will be added between the tracker and ECAL that provides a precise measurement of when each particle passed through it. This effectively adds another dimension (time) to particle identification, allowing for greater mitigation of pileup.
- *High granularity calorimeter* [99, 100]: The forward sections of the ECAL and HCAL will be entirely replaced with a new calorimeter that incorporates layers of silicon detectors. This will give timing information for PU mitigation, higher resolution, and more.

The muon chambers and trigger system will also see significant upgrades, along with a major overhaul of much of the less glorious hardware.

2.4 Acknowledgements

The figures used in this chapter are materials produced by or for CERN and were released by CERN for informational use under CERN copyright [101]. The descriptions of the LHC and CMS were derived from publicly available material and the relevant technical design reports.

Chapter 3

Studying the Higgs boson through vector boson scattering at CMS

It's very nice to be right sometimes.

(Peter Higgs, Edinburgh University press conference, 2012)

3.1 New physics and the Higgs boson

Chapters 4 and 5 cover two distinct analyses of CMS data—hereafter, we will use the word “analysis” as a shorthand for any publishable study of CMS data. Both chapters report precision measurements of the Higgs boson as an exciting handle for discovery via similar physics processes. As such, in this chapter, we will introduce common strategies, terminology, and techniques used in both analyses. Then, in the following chapters, we will narrow in on specifics.

The stage is set: over a century of particle physics has yielded a beautiful theory of *almost* everything, the Standard Model, and a grand coalition of nations has built the largest and most complex scientific instrument in human history, the LHC, to test it. The most recent act in this drama concluded in 2012, when the Higgs boson was discovered at CMS and ATLAS [24, 23]. In the years following its discovery, the LHC experiments have measured many of its properties to great precision and found no significant deviations from SM predictions [102, 103]. However, there are many features yet to be measured precisely. This implies that studies of the Higgs boson

have incredible potential for discovery: a better understanding of the Higgs boson immediately yields a better understanding of the origin, operation, and continued existence of the entire universe. Moreover, the Higgs boson can uniquely interact with all known matter, so by producing many Higgs bosons in the lab, we may be able to access physics beyond the Standard Model (BSM).

There are many educated guesses, called theories, as to the mysterious nature of the Higgs boson. Some guess at the existence of yet-undiscovered particles that also interact with the Higgs boson¹, and experimentalists can search for their existence directly: either by looking for the SM particles that they might decay into, or by checking for what is missing after everything is else accounted for. Experimentalists can also search for new physics indirectly by making precise measurements of SM predictions: any significant deviation from the prediction would poke another hole in the Standard Model or even confirm a prediction of a new theory. The physics analyses described in this document both follow the latter strategy.

3.1.1 The κ -framework

One commonly used framework used to quantify deviations from the SM is the so-called κ -framework [104], which defines modifiers κ_X to the Higgs boson couplings to any particle X :

$$\kappa_X = \frac{\text{modified coupling value}}{\text{SM coupling value}}. \quad (3.1)$$

While there are myriad theoretical nuances to the statement above, it is sufficient to state the obvious: $\kappa_X = 1$ represents the SM scenario and deviations from 1 represent BSM scenarios. The κ -framework is not the only framework used to understand and quantify modifications to the SM, however, with the most notable alternative being Effective Field Theory [105, 106].

¹This is well-motivated: dark matter is known to have mass, and it may have obtained that mass through the Higgs mechanism.

3.2 Producing Higgs bosons with vector boson scattering

The dominant Higgs boson production mechanism at the LHC is gluon-gluon fusion through a top quark loop. Through these processes, the Higgs boson was discovered and many of its branching ratios and couplings have been measured precisely. However, additional production mechanisms can yield unique insights. Enter vector boson scattering² (VBS), where two quarks scatter off of each other by the exchange of vector bosons ($V = W$ or Z), as shown in Fig. 3.1. This exchange can, in turn, produce a variety of physics, most notably the production of at least one Higgs boson. This process also has two important features. First, it has a unique signature: the two scattered quarks fly out of the collision back-to-back, typically in the forward regions of the detector (Fig. 3.6), with large energies. Second, but equally important, modifications to the Higgs couplings induce effects that grow with energy [108]—at the LHC, where energy is plentiful, this makes BSM scenarios easier to spot.

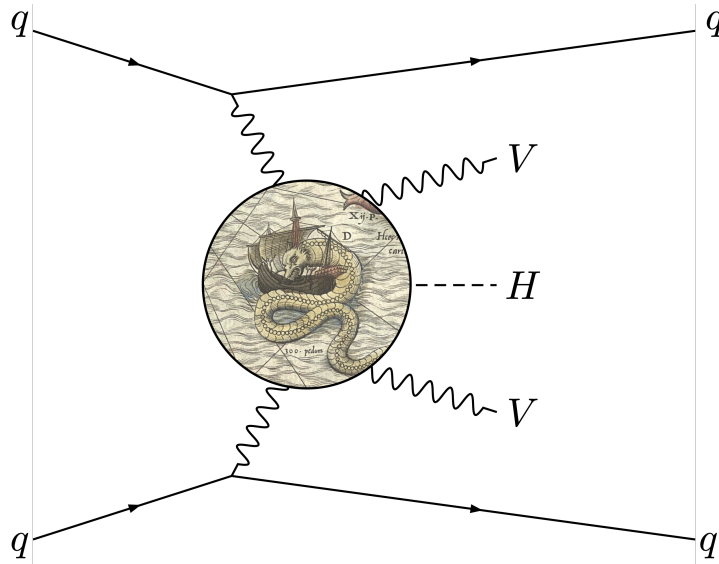


Figure 3.1. The leading-order Feynman diagram for VBS production of a Higgs boson and two vector bosons. This can occur in several ways, represented by a circle in the middle obscuring the vertices involved (“here be dragons”). Inspired by Ref. [109].

²Although VBS specifically implies the production of two vector bosons, while vector boson fusion (VBF) implies the production of a single Higgs or vector boson [107], the names VBS and VBF are used interchangeably; since we are interested in the production of a Higgs boson and one or two vector bosons, the naming is anyway ambiguous.

VBS provides a unique way of producing Higgs bosons, with an easily identifiable signature, and there are myriad VBS processes involving a Higgs, each involving different couplings. The rarer couplings are of particular interest, including the HHVV and HHH couplings (Ch. 5). Also of interest are some of the more obscure properties of the Higgs boson, like the relative sign between the HWW and HZZ couplings (Ch. 4). In all of these measurements, if we confirm SM predictions, we exclude possible BSM directions, and if we find a significant deviation from the SM, we will have opened an entirely new chapter of physics.

We begin by selecting the Feynman diagrams that involve the couplings of interest—here, we focus on VBS processes. We call the selected processes our “signal.” For example, if we are interested in HHVV, as in Ch. 5, our signal might be VBS production of a Higgs boson and two vector bosons. We must then find our signal, which is often incredibly rare, amongst petabytes of data recorded by CMS.

3.3 Choosing a final state

The final collection of particles that are actually recorded by CMS are called the “final state.” As covered in Chapter 2, Section 2.2, CMS can record electrons, photons, hadrons, and muons. Higgs bosons, vector bosons, gluons, and tau leptons decay to these particles. Therefore, when we say we are looking for VBS production of a Higgs boson and two vector bosons, we are really looking for two quarks (VBS) and the decay products of the bosons. Some final states are preferable over others, however. For convenience, we often prefer to look for $H \rightarrow b\bar{b}$, since it occurs 58% of the time—the next-largest “branching ratio” is $H \rightarrow W^+W^-$ at 21%. The choice for the decay of the vector bosons is less clear. The leptonic decays ($W \rightarrow \ell\nu$ or $Z \rightarrow \ell^+\ell^-$) are more unique, but are rarer, with branching ratios of 33% and 10%, respectively. On the other hand, the hadronic decays ($W \rightarrow qq$ at 67% and $Z \rightarrow qq$ at 70%) are more plentiful, but are difficult to distinguish from far more abundant processes. The *reconstruction* of the final state of each proton-proton collision, based on the electronic readout of CMS, is thus crucial for the

entire physics operation at CMS—and any other experiment at the LHC. Importantly, many SM particles decay to neutrinos, which pass through CMS completely undetected, and it is possible that BSM particles are similarly missed, but the presence of these particles can be inferred if the reconstruction is performed with high precision (see Section 3.4.3).

3.4 Reconstruction

The subdetectors that comprise CMS produce a set of electric signals digitized into data that can be used as measurements of physical quantities. Moreover, when associated to one another, these measurements can also be used to ascertain the identity of detectable throughgoing particles (Fig. 3.2). This is done via a global “particle-flow” (PF) algorithm [110], which combines the information from each of the CMS subdetectors to reconstruct all individual particles in an event. Particles reconstructed by the PF algorithm are called PF “candidates” and are used to build jets, which originate from hadrons, τ leptons, and in the computation of the missing transverse momentum (p_T^{miss}).

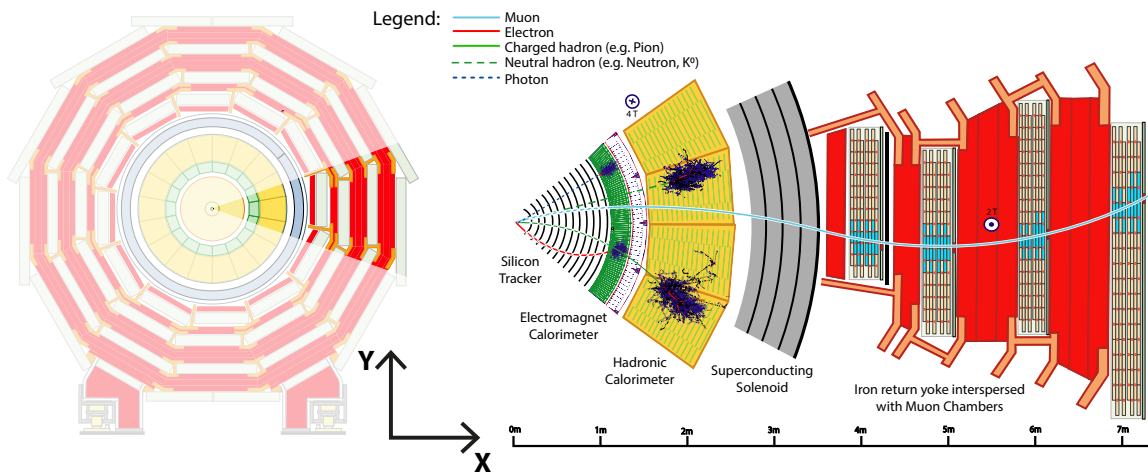


Figure 3.2. A transverse slice of CMS, from Ref. [111], showing the signals left by different kinds of particles in each subdetector.

3.4.1 Jets

All hadrons leave a deposit in the HCAL, while charged hadrons also have a track that can be associated to it, but these hits do not correspond to individual quarks—like the VBS quarks, for example. Due to “color confinement,” quarks hadronize into “jets” of hadrons. This happens in a cascade, where quarks beget quarks that beget more quarks and so on, until the jet impinges on the calorimeters (Fig. 3.3). Fortunately, jets tend to be concentrated in relatively narrow cones, meaning individual jets—and thus, individual quarks—can usually be distinguished from one another. Jets are constructed using the anti- k_T algorithm [112, 113], which clusters PF candidates according to a configurable “distance parameter,” which determines the size of the jet cone. For AK4 jets, the distance parameter is set to 0.4, corresponding to $\Delta R = 0.4$, where

$$\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2} \quad (3.2)$$

and $\Delta\phi$ and $\Delta\eta$ are the difference in azimuthal angle and pseudorapidity (defined in Appendix A) between two vectors, respectively. There is also a collection of AK8 jets, which have the distance parameter set to 0.8, for reconstructing merged jets.




Figure 3.3. Illustration of two quarks hadronizing into two distinct jets based on Figs. 8.1, 8.2, and 8.3 in Ref. [3].

Jets originating from b quarks

While the presence of quarks can be inferred from the presence of jets, the identity (flavor) of each quark is less obvious. However, when b quarks hadronize, they produce a B meson (Fig. 3.4a). The B meson lifetime is sufficiently long such that it passes the first few layers of the tracker before decaying—only picoseconds after it was produced. When it decays, it produces several quarks and, potentially, a lepton. This results in “displaced” tracks in the tracker that point to a “secondary vertex” (Fig. 3.4b) as opposed to the “primary vertex,” or the original pp collision. The presence of a secondary vertex within a jet therefore gives a unique indication that it originated from a b quark (Fig. 3.4c). However, the reconstruction of that secondary vertex is fundamentally made possible by the fact that the B meson is sufficiently more massive than its decay products. If the B meson were much lighter, its decay products would be more collimated, and the secondary vertex could therefore only be reconstructed if the tracker were infinitely precise. Particle tracks are therefore vital for b-tagging, so b-tagged jets must be within the tracker acceptance region ($|\eta| < 2.5$). For CMS analysis, a deep neural network called DEEPJET was trained to use the presence of a secondary vertex, along with additional kinematic information about the jet and its constituents, to “tag” jets as originating from a b quark [114].

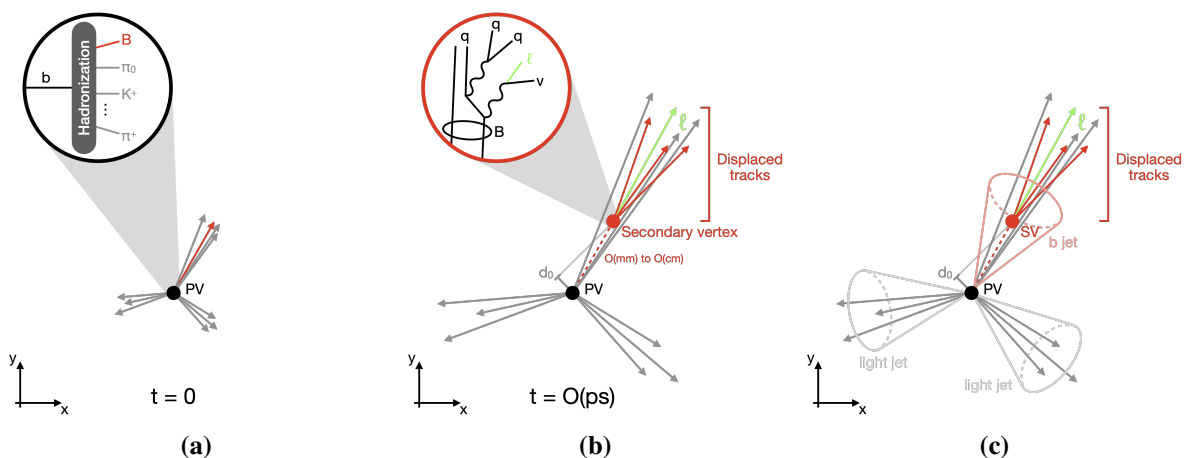


Figure 3.4. A sketch of the hadronization of a b quark (left), the subsequent hadronization of the produced B meson (center), and the identification of each jet (right). Based on Ref. [115].

Merged jets originating from $H \rightarrow b\bar{b}$

In cases where a highly energetic particle decays to a pair of quarks, the resulting jets overlap or are near enough to each other to be difficult to separate (Fig. 3.5). This presents a unique signal: the two nearby jets can instead be found as a single, large-cone jet. That is, these merged jets will appear in the AK8 jet collection. In the work presented in this dissertation, we search for Higgs bosons decaying to $b\bar{b}$ at high energies, tagged as a single merged jet. This is made possible by a graph neural network, called PARTICLENET, trained for CMS analysis to identify the origins of merged jets based on the properties of its constituents, arranged in a graph-like data structure according to their proximity to one another [116]. Similar to b-tagging, merged jets considered for $X \rightarrow b\bar{b}$ tagging, for instance, must be within the tracker acceptance region. The merged jet with a mass similar to the Higgs boson, and the highest PARTICLENET $X \rightarrow b\bar{b}$ score is considered as the best Higgs boson candidate.




Figure 3.5. Illustration of a generic boson with a small, moderate, and large Lorentz boost, respectively, decaying to two jets. As the boson receives a larger boost, its hadronic decay products become increasingly collimated until they are reconstructed as a single merged jet. Based on Ref. [117].

Jets originating from VBS quarks

The presence of VBS quarks is inferred from the presence of two jets that are back-to-back (large $|\Delta\eta_{jj}|$) and have a large combined invariant mass (M_{jj}). In events where there are more than two additional jets that do not overlap with the Higgs boson candidate jet, the two

best VBS jet candidates must be selected. A unique VBS selection technique was derived for the work presented in this dissertation: if all of the jets are in one hemisphere of η , then the two jets with the largest momenta are selected; otherwise, the jet with the largest momentum in each η hemisphere are selected.

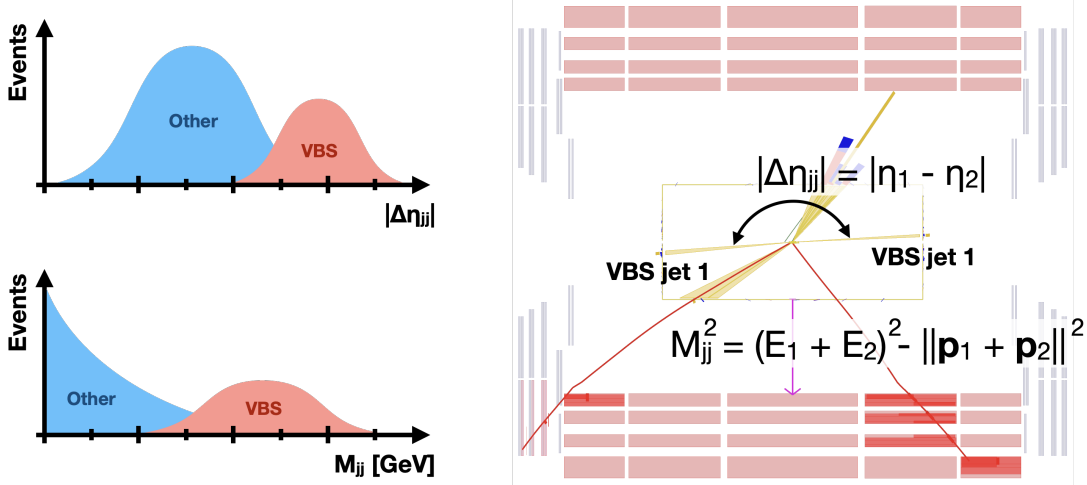


Figure 3.6. Typical VBS jet M_{jj} and $|\Delta\eta_{jj}|$ distributions (left) next to an event display for a simulated VBS WH signal event in the r - z plane, with the VBS jets labeled (right). M_{jj} is the invariant mass of the two VBS jets, and $|\Delta\eta_{jj}|$ is the pseudorapidity (defined in Appendix A) gap between them.

3.4.2 Leptons

For simplicity, when we consider “leptons,” we will exclusively mean those leptons that are directly detected by CMS, namely electrons and muons. Electrons have a track in the silicon tracker and leave a deposit in the ECAL. The tracks and ECAL deposit can be associated to one another, distinguishing electrons from photons, which only leave a deposit in the ECAL. Muons also have a track, but they fly past both calorimeters undetected and leave hits behind in the muon chambers.

In the SM, the lepton family also includes tau leptons and neutrinos. Tau leptons decay almost immediately to hadrons or to a tau neutrino and W boson. The latter decay is difficult to reconstruct, because it is easily confused with a lone W boson. The hadronic decay, however,

has a unique “multi-prong” structure, allowing it to be distinguished from other jets. Neutrinos, meanwhile, pass through the entire detector, so their presence can only be inferred by first accounting for everything else, then asking for what is missing.

3.4.3 Missing transverse momentum

Protons collide along the beamline with equal but opposite momenta, so the system has a total momentum of zero. Because of the conservation of momentum, the sum of the momenta of all produced particles must still be zero in the transverse plane. Importantly, the interacting protons receive a random fraction of the proton’s overall momentum according to the parton distribution function, so the initial momentum of the produced particles in the z -direction is not known a priori. The picture in the transverse plane can meanwhile be visualized in the classical world. The shrapnel from two colliding bullets, for instance, is thrown from the collision in a disk such that the total momentum is conserved [118]. Therefore, if an invisible particle (a neutrino or possibly something new) is produced, its momentum will not be recorded, and the visible p_T sum will thereby be non-zero. The “missing” p_T required to conserve transverse momentum is referred to as p_T^{miss} .

3.5 Identifying backgrounds

Of course, the observed final state is not unique. For example, if we were searching for $Z \rightarrow \ell^+ \ell^-$, and thus asked for collision events with two leptons with opposite charge, we would also get events where a W^+ and W^- were produced and both decayed leptonically. We could also pick up events where W and Z boson were produced, and both decayed leptonically, but the lepton from the W was not recorded by CMS. These processes that produce the same final state, but are not the signal, are referred to as “backgrounds.”

One common background for VBS Higgs boson analyses is $t\bar{t}$ production. Both top quarks decay to a W boson and b quark. With either the hadronic or leptonic decay of the W bosons, and the two genuine b jets, there is a large potential for trickery. For example, the W

boson could decay to two quarks that overlap with each other and one of the b jets, resulting in a jet with a large invariant mass and displaced tracks that could be mistaken for a $H \rightarrow b\bar{b}$ jet (e.g. Fig. 3.7). Backgrounds including genuine VBS jets or a genuine Higgs boson are typically much rarer than a more populous background like $t\bar{t}$ generating a fake signature.




Figure 3.7. Event display for a simulated VBS WH signal event (left) and $t\bar{t}$ event (right). Both events contain two forward jets, a single muon, and two b jets.

3.6 Simulation

Some experiments can control when interesting data is being recorded, so background noise can be directly measured by simply turning the source of interesting data off, but leaving the experiment on. The CMS Experiment does not have this luxury: interesting collision events are buried amidst billions of uninteresting ones—the L1 trigger and HLT only make very general selections. Moreover, as covered in the previous section, the signal final state is not unique, so although the PF algorithm is capable of reconstructing each particle produced in a proton-proton collision, we must carefully examine those produced by signal versus background such that we may separate them in our analysis. This cannot be done easily using only data, since we do not know a priori what really happened—we can only know what was produced. Instead, we use Monte Carlo (MC) simulation software that “generates” proton-proton collisions, where we can

specify exactly what physics processes occur [119].

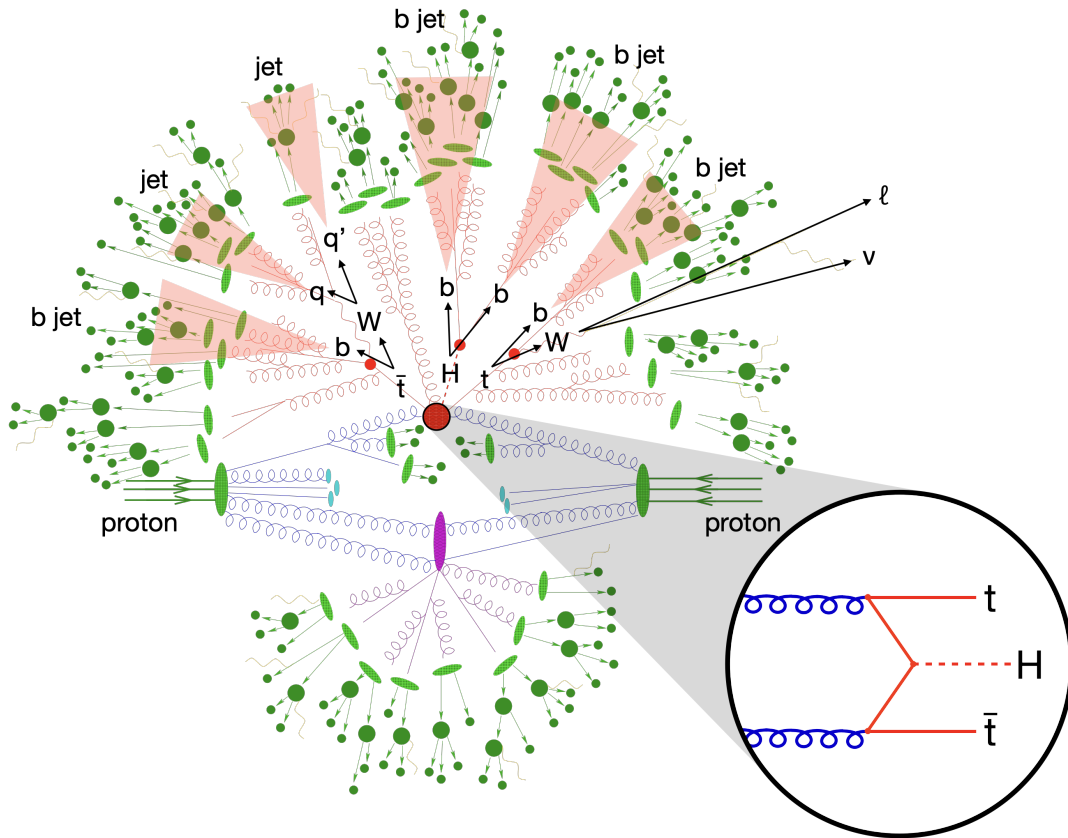


Figure 3.8. Detailed diagram of the generation of a single $pp \rightarrow t\bar{t} + H$ event, based on Ref. [120], starting from hard scattering (circled), to the decay of the top quark, anti-top quark, and Higgs boson (black arrows and text), to the hadronization of the final state into jets (red triangles). Also included are jets from final-state radiation, i.e. gluons radiated from the final state quarks, and initial-state radiation, i.e. jets from the other constituents of the colliding protons.

Taking $pp \rightarrow t\bar{t} + H$ as an example (Fig. 3.8), the first step of generating simulated collision events is to calculate the kinematics (the energy and the momentum vector) for the top quark, anti-top quark, and Higgs boson. This is often done with the MADGRAPH5_aMC@NLO “matrix-element” generator [121], which finds the relevant Feynman diagrams, including very rare ones, and randomly generates different realizations of the final state based on an approximation of the underlying theory. Better accuracy comes from including more Feynman diagrams and a more fine-grained approximation of QFT at the cost of computational complexity. Next, a “general-purpose” generator (e.g. PYTHIA or POWHEG [122, 123, 124]) takes the event generated by the

previous step³ and promotes it to something closer to what we would observe in nature. That is, it generates the particles that actually impinge onto our detector, including an approximation of the hadronization of quarks and gluons into jets. The output of this step is then interfaced with a highly detailed GEANT4 simulation [125] of the CMS Experiment—from material interactions, to the digitization of the electrical signals from the subdetectors, to the decisions of the L1 trigger and HLT. The result is a sample of simulated data that is a very close match to the signals produced by CMS for $pp \rightarrow t\bar{t} + H$.

Because it is simulated all of the way to the trigger decisions, the simulation can be passed through the PF algorithm and all other data post-processing steps performed on data (Fig. 3.9). This results in samples of simulated data that appear effectively identical to real collision data, except with some truth-level information available for further study. Analyzers can then use these simulated samples to optimize their analysis for preferring signal processes over background processes.

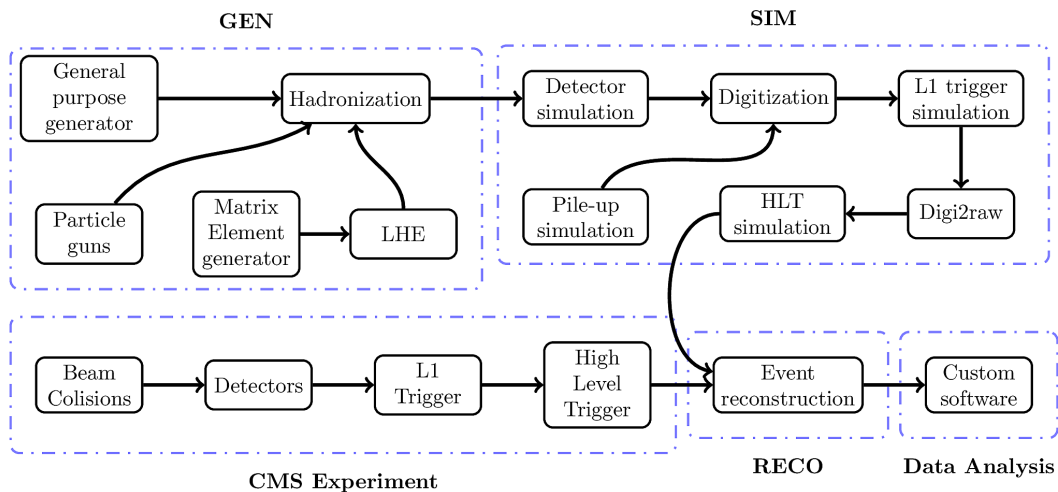


Figure 3.9. Diagram of the CMS collision data and simulation processing workflow, from [126], for generating MC simulation showing the post-processing steps that are shared between MC simulation and real collision data.

³Notably, general-purpose generators are also capable of doing the first step.

3.7 Event selection

In order to isolate the signal, we must place a set of selections on our data that prefer events that look like signal. We do this in a particular order, however, because of the sheer size of our data: over the course of Run 2, the CMS Experiment collected over 200 PB of data! This data is processed and re-processed into more and more condensed descriptions of the same recorded events—the simulation is identically re-processed—but even at the smallest so-called “data tier”, there is 200 TB of data and 900 TB of simulation (Fig. 3.10). Starting with triggers, progressively strict selections are placed that are optimized to prefer events that contain our signal. This process, orchestrated across 170 computing clusters in over 40 countries [127] (i.e. 1000s of distributed CPUs), reduces 1.1 PB of data down to a few histograms.

We first select the HLT paths most relevant to the chosen final state. For example, if the final state includes one lepton—say, from the leptonic decay of a W boson—then then we would select the single-lepton HLT paths. This isolates a portion of the petabytes of CMS data that are relevant for analysis. The HLT paths are also implemented in the simulation, so we will require that simulated events pass the same HLT paths as the data.

Next, a series of selections are placed on the events passing the HLTs in order to narrow in on a set of events most relevant for analysis. For VBS Higgs boson analyses, these will select for at least two jets that look like VBS jets, and $H \rightarrow b\bar{b}$ reconstructed as one merged jet or two separate jets. These selections are referred to as the “Preselection” in this dissertation.

The final set of selections are optimized using events that pass the Preselection. In the simplest analyses, like that presented in Chapter 4, these selections will be on properties of each particle in the signal final state. Other analyses, like that presented in Chapter 5, use these variables as input to a machine learning algorithm trained to select signal events. Ultimately, we select a final set of events that comprise the so-called “signal region” (Fig. 3.11). We will compare the number of actual collision events that enter this region to the number of events predicted by simulation. However, we must first derive the best possible estimate of the background in this

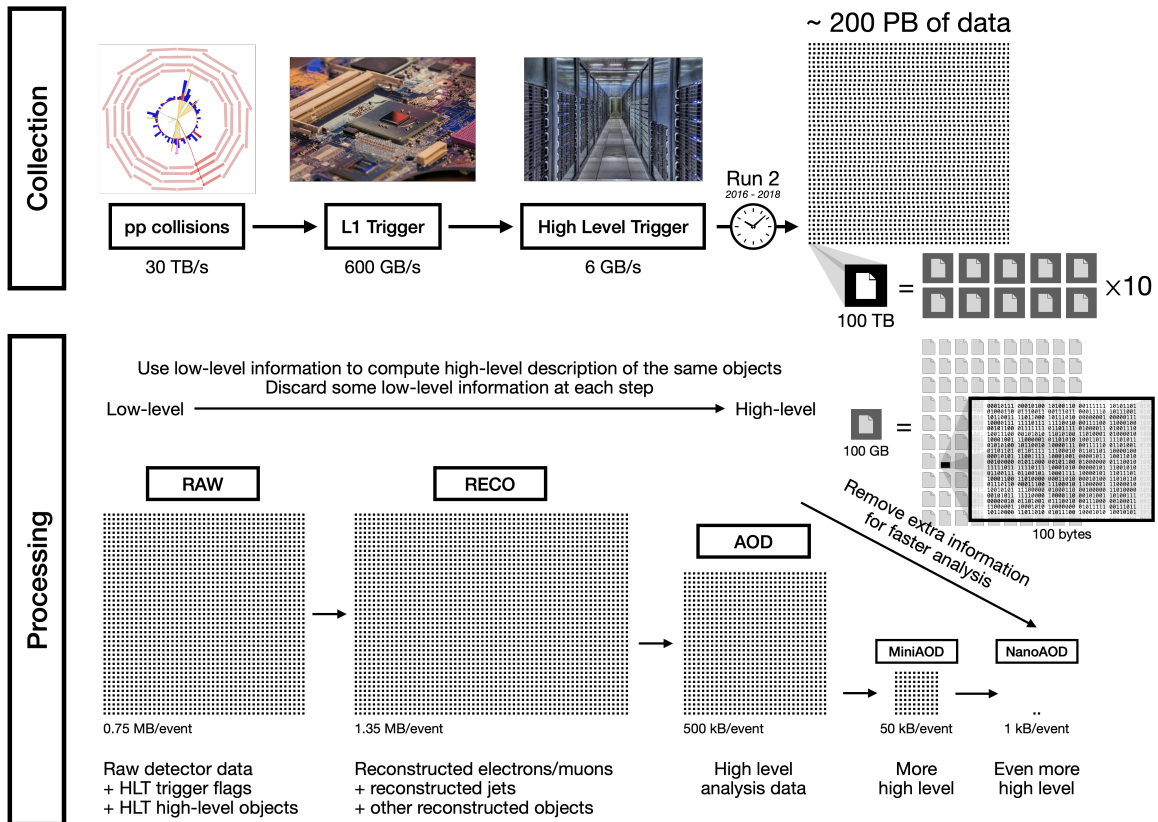


Figure 3.10. Diagram of the evolution of CMS data from collection (top) to post-processing into the data used in analysis (bottom). The 30 TB/s of data that are produced by the 40 MHz proton-proton collisions is reduced to 6 GB/s by the trigger system, which discards uninteresting events. Over the course of Run 2, this yielded over 200 PB of recorded data that was then processed into different data “tiers” for further analysis.

region and compute all relevant systematic uncertainties.

3.8 Background estimation

Although the predicted number of background events in the signal region can be taken directly from the simulation, the simulation is not perfect. Instead, there is a strategy—the famous “ABCD” method—in which only real data is used. The ABCD method entirely relies on there being two variables used to define the signal region that are decorrelated. For example, an analysis may have variables x and y , where most of the signal has $x > 4$ and $y > 200$, whereas the background is concentrated near small values of x and y (Fig. 3.12). In this case, if x and y are

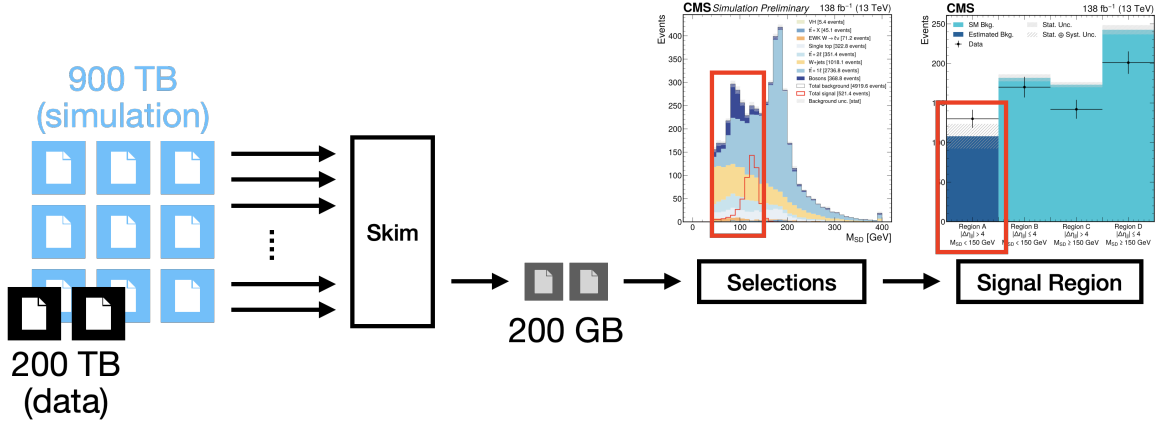


Figure 3.11. Diagram of the data processing workflow for a typical measurement with CMS data. Starting from NanoAOD (200 PB of real collision data, 900 PB of simulation), “skimming” jobs are distributed across thousands of GPUs, each taking a few NanoAOD files (around 1 GB each) and removing all of the unnecessary information. Then, progressively strict selections are placed until the signal region is reached, wherein the measurement will be made.

decorrelated, then for any value of x , the distribution of the background in y will be the same and vice versa. Provided that this is true, we can look at real collision events with $x \leq 4$ and calculate

$$f = \frac{\text{N events with } y > 200 \text{ and } x \leq 4}{\text{N events with } y \leq 200 \text{ and } x \leq 4} \quad (3.3)$$

The distribution of y , and thereby the fraction f , should be identical to that for events with $x > 4$, since x and y are decorrelated, i.e.

$$f' = \frac{\text{N events with } y > 200 \text{ and } x > 4}{\text{N events with } y \leq 200 \text{ and } x > 4} \quad (3.4)$$

$$f' \approx f \quad (3.5)$$

Therefore, we can take the number of events with $x > 4$ and $y \leq 200$ and multiply it by f to get the number of events we expect to have $x > 4$ and $y > 200$, i.e.

$$\left(\text{N events with } y > 200 \text{ and } x > 4 \right) \approx \left(\text{N events with } y \leq 200 \text{ and } x > 4 \right) \times \frac{\text{N events with } y > 200 \text{ and } x \leq 4}{\text{N events with } y \leq 200 \text{ and } x \leq 4} \quad (3.6)$$

For brevity, let A , B , C , and D , be the number of events with $y > 200$ and $x > 4$, $y \leq 200$ and $x > 4$, $y > 200$ and $x \leq 4$, and $y \leq 200$ and $x \leq 4$, respectively. Then, the equation above becomes

$$A \approx B \times \frac{C}{D} \quad (3.7)$$

By estimating the background yield in the signal region in this way, we rely entirely on nature, which has no associated systematic uncertainties.




Figure 3.12. A generic ABCD configuration where selections on the statistically independent quantities x and y are used to form the signal region (region A) and three “control” regions (regions B, C, and D) used to estimate the background (black dots) in the signal region. In the toy example shown, the method closes exactly: $B \times C/D = 10 \times 6/15 = 4$.

3.9 Systematic uncertainties

Since the background estimation strategy is data-driven, the systematic uncertainties on the simulation, which are numerous, only need to be evaluated for the signal yield. Systematic uncertainties in CMS statistical analysis come in two flavors: theoretical and experimental uncertainties. Sources of experimental uncertainty include the efficiency of the HLT, lepton identification, and b tagging as well as the scale and resolution of the jet energies. For each of these sources, a correction has been derived to make simulation behave more like data. The HLT efficiency, for example, can be directly compared to data; 90% pass in simulation, versus

95% in data, then each passing event can be considered instead as 1.056 events, forcing the overall yields to match. Then, the associated experimental uncertainty is quantified by taking the correction, varying it up and down by one standard deviation⁴, and assessing the impact on the final measurement. Theoretical uncertainties, meanwhile, are quantified by varying parameters of the simulation up and down and assessing the impact on the final measurement.

3.10 Statistical interpretation

The analyses presented in this dissertation are counting experiments. That is, for the final result, we take the expected signal yield A_{signal} and predicted background yield $A_{\text{data}}^{\text{pred}}$ in the signal region and compare them to the yield of real collision events A_{data} . Intuitively, if $A_{\text{data}} = A_{\text{data}}^{\text{pred}}$, then our signal most likely does not exist in nature; conversely, if $A_{\text{data}} = A_{\text{signal}} + A_{\text{data}}^{\text{pred}}$, then we have just confirmed that our signal does exist. But what can we say if A_{data} is somewhere between $A_{\text{data}}^{\text{pred}}$ and $A_{\text{signal}} + A_{\text{data}}^{\text{pred}}$? We have also just enumerated through the many systematic uncertainties on A_{signal} and $A_{\text{data}}^{\text{pred}}$ that affect our ability to make this comparison exactly. For example, $A_{\text{data}}^{\text{pred}}$ may have, by chance, fluctuated upward by just enough to make it seem as though our signal exists—or worse, it may have fluctuated downward and made us think our signal is not real. Using statistics, we can make a more precise statement about our measurement, taking all of the uncertainties into account.

3.10.1 Probability and likelihood

The *probability* that we observe n events in our signal region given the expected yield λ follows Poisson statistics:

$$P(n|\lambda) = \frac{\lambda^n e^{-\lambda}}{n!} \quad (3.8)$$

In our analysis, $\lambda = \mu s + b$, or some fixed number μ times the signal yield s plus the background yield b . The number μ is called the “signal strength” because μs represents different scenarios:

⁴This uncertainty is derived differently for each correction.

e.g. the signal does not exist ($\mu = 0$), the signal does exist ($\mu = 1$), the signal is 50% smaller than our model expected ($\mu = 0.5$), and so on. Revising our equation, we obtain

$$P(n|\mu, s, b) = \frac{(\mu s + b)^n e^{-(\mu s + b)}}{n!} \quad (3.9)$$

We need a slightly different formulation, however, because we have already observed n events. That is, we already know the *outcome* of our experiment, and we are trying to explain that outcome based on the *parameters* of the underlying statistical model—Poisson statistics in this case. In other words, we need the probability that the signal strength is some number μ , given the yields n , s , and b . We call this the likelihood L , or

$$L(\mu) = \frac{(\mu s + b)^n e^{-(\mu s + b)}}{n!} \quad (3.10)$$

Importantly, P and L are related, and thus superficially similar, but they are distinct quantities (Fig. 3.13).




Figure 3.13. The probability for different values of μ (left) and the likelihood for $n = 1$ (right), showing how $L(\mu)|_{n=1}$ is given by computing $P(1|\mu)$.

3.10.2 Maximum likelihood estimate

Our task can now be expressed in the proper language of statistics: determine the most probable value of μ given n , s , and b by calculating the maximum of the likelihood $L(\mu)$.

Fortunately, this is basic calculus: take the derivative of L with respect to μ , set it equal to zero, and solve. For simplicity, we will instead consider $q(\mu) = -2\ln L(\mu)$:

$$\begin{aligned}
 q(\mu) &= -2[n\ln(\mu s + b) - (\mu s + b) - \ln(n!)] \\
 \frac{\partial q}{\partial \mu} &= 1 - \frac{n}{\mu s + b} \quad \text{where } \left. \frac{\partial q}{\partial \mu} \right|_{\mu=\hat{\mu}} \equiv 0 \\
 0 &= 1 - \frac{n}{\hat{\mu} s + b} \\
 \hat{\mu} &= \frac{n - b}{s}
 \end{aligned} \tag{3.11}$$

In our simple setup, this is rather intuitive: plugging $\hat{\mu}$ into $\mu s + b$ gives n exactly. We can now easily compute $\hat{\mu}$ for a given value of n , s , and b , but we must also declare how confident we are in our measurement.

Assuming the likelihood is approximately Gaussian⁵ with an amplitude \mathcal{N} and standard deviation of σ_μ —which is true for large n (Fig. 3.14)—we can express $L(\mu)$ as

$$L(\mu) \approx \mathcal{N} e^{-\frac{(\mu - \hat{\mu})^2}{2\sigma_\mu^2}} \tag{3.12}$$

For convenience, we again consider a more useful quantity $q(\mu)$:

$$q(\mu) = -2\ln \frac{L(\mu)}{L(\hat{\mu})} = \left(\frac{\mu - \hat{\mu}}{\sigma_\mu} \right)^2 \tag{3.13}$$

where $\sqrt{q(\mu)}$ can be interpreted as the number of standard deviations away μ is from $\hat{\mu}$. In other words, for a given value of $\hat{\mu}$, we can now express the range of possible values of μ within 1, 2, or more standard deviations (Fig. 3.15)—this is called a confidence interval. If, for example, we had 413 expected signal events, 128 predicted background events, and observed 550 real events, then we could say that the signal strength was measured to be within the 95% confidence

⁵For non-physicists, this is referring to a normal distribution or bell curve.

interval $[0.88, 1.10]$.




Figure 3.14. The likelihood with the Gaussian approximation of the likelihood overlaid for different values of n , showing poor agreement for small n (left), decent agreement for $n = 56$, and good agreement for $n = 550$.


So far, we have not considered the uncertainty in s or b , which are critical to the accuracy of our confidence interval. Focusing on the signal yield, we can decompose s into multiple components:

$$s = \underbrace{\left(\text{detector acceptance} \right)}_A \times \underbrace{\left(\text{detection efficiency} \right)}_\varepsilon \times \underbrace{\left(\text{integrated luminosity} \right)}_{\mathcal{L}} \times \underbrace{\left(\text{cross section} \right)}_\sigma \quad (3.14)$$



At CMS, we are able to measure the integrated luminosity for all of Run 2 with an uncertainty of around 1.6%. If we wanted to include this uncertainty in our statistical analysis, we could rewrite the luminosity term \mathcal{L}

$$s = A\varepsilon\mathcal{L}_0(1 + 1.016)^\eta\sigma \quad (3.15)$$

where \mathcal{L}_0 is the measured value of the integrated luminosity and η is typically a Gaussian-distributed number between 0 and 1. Before we compute our confidence interval, we must find some $\hat{\eta}(\mu)$ that minimizes $q(\mu, v(\mu))$. For each systematic uncertainty in our analysis, we add another “nuisance parameter” that must first be minimized before the final result is produced. In this dissertation, this operation—along with the maximum likelihood estimation—is performed by the COMBINE tool [128].



signal exists!



signal does not exist!

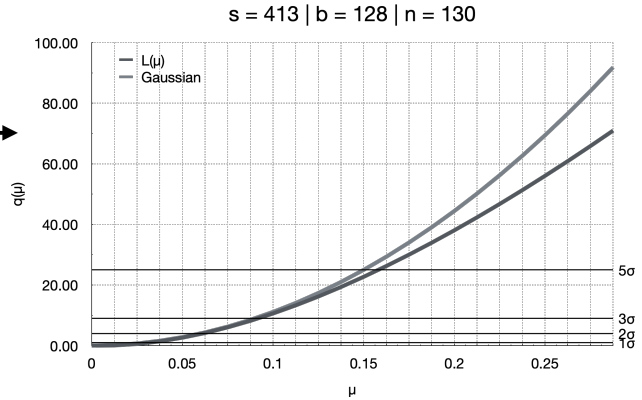


Figure 3.15. The likelihood (left) and confidence interval (right) plotted for two different scenarios: $n \approx s$, or the signal exists, (top) and $n \approx b$, or the signal does not exist (bottom). In the first scenario, the signal strength is measured to be within $[0.88, 1.10]$ at the 95% CL, whereas in the other scenario, it is measured to be within $[0, 0.06]$ at the 95% CL.

Chapter 4

Determination of λ_{WZ} through VBS WH production

4.1 Looking for a sign

Before this work was published, the CMS Collaboration had measured the *magnitudes* of the HWW (κ_{W}) and HZZ (κ_{Z}) couplings to be $|\kappa_{\text{W}}| = 1.02^{+0.11}_{-0.10}$ and $|\kappa_{\text{Z}}| = 1.04^{+0.07}_{-0.07}$, showing precise agreement with the SM values [102]—with similarly strong constraints from ATLAS [103]. The *sign* of either coupling, however, had not yet been well-determined. The relative sign between κ_{W} and κ_{Z} , which was of particular interest, can be expressed more compactly as the ratio between the two couplings:

$$\lambda_{\text{WZ}} = \frac{\kappa_{\text{W}}}{\kappa_{\text{Z}}}. \quad (4.1)$$

The Standard Model requires $\lambda_{\text{WZ}} = 1$ in order to preserve the “custodial” symmetry. Meanwhile, certain BSM theories require λ_{WZ} to be negative [129], including models with higher isospin representations [130] like the Georgi-Machacek model [131]. Therefore, in the absence of a significant experimental measurement of the sign of λ_{WZ} , a crucial element of the SM had not been confirmed, and a potential signature of BSM physics was left unexplored.

4.2 The signal

The precise determinations of the magnitudes of κ_W and κ_Z [102] were obtained in studies of processes that are predominantly quadratic in κ_V —that is, κ_W or κ_Z enter the Feynman diagram twice. While there were some with a linear dependence, they did not give a strong exclusion of opposite-sign scenarios [132]—in fact they slightly preferred the $\lambda_{WZ} < 0$ scenario. A promising channel to directly probe λ_{WZ} at the LHC is the production of VH via vector-boson scattering (VBS) [133], shown (at leading order) in Fig. 4.1. Such a channel is sensitive to the relative sign of κ_W and κ_Z since the cross-section σ has an interference term that is linear in both κ_W and κ_Z [133]:

$$\sigma \propto |\mathcal{M}|^2 = \kappa_W^2 |\mathcal{M}_W|^2 + \kappa_W \kappa_Z \mathcal{M}_{WZ}^2 + \kappa_Z^2 |\mathcal{M}_Z|^2 \quad (4.2)$$

where the matrix elements for the contributions from the HWW couplings, HZZ couplings, and interference term are denoted as \mathcal{M}_W , \mathcal{M}_Z , and \mathcal{M}_{WZ} , respectively. Therefore, this channel provides the opportunity to determine the sign of λ_{WZ} .




Figure 4.1. Leading-order Feynman diagrams for VBS production of a W and Higgs boson, where the W decays leptonically and the Higgs decays to b quarks. The Higgs coupling to W bosons κ_W and Z bosons κ_Z is denoted by a red circle (●).

4.2.1 Signal characteristics

A specific final state is intentionally selected for its uniqueness—making the signal easier to find amidst the haystack. First, leptons are preferred in the final state over jets, due to the

sheer size of the ‘‘QCD multijet’’ background, the most populous physics process produced at the LHC where the interactions between the quarks and gluons in the colliding protons simply produce more quarks and gluons. Then, $W \rightarrow \ell\nu$ is preferred over $Z \rightarrow \ell\ell$, since there are fewer backgrounds with only one lepton in the final state. Finally, $H \rightarrow b\bar{b}$ has the largest branching ratio, and it can be identified using the latest advances in artificial intelligence.

While there are a few, quite large background processes can produce the target final state, the Monte Carlo (MC) simulation of the signal reveals additional boons. The final state particles receive a significant boost for negative λ_{WZ} scenarios (Fig. 4.2). In our preferred final state, this will give a lepton with large p_T , some additional E_T^{miss} from the neutrino, and two overlapping b-jets reconstructed as a single merged jet. This can be quantified in a single variable by simply summing the p_T of each final-state particle other than the VBS jets:

$$S_T = p_T(W \rightarrow \ell\nu) + p_T(H \rightarrow b\bar{b}) = p_T(\ell) + p_T^{\text{miss}} + p_T(\text{AK8 jet}) \quad (4.3)$$

The cross-section of the signal process increases almost quadratically as κ_W or κ_Z deviate from the SM value (Fig. 4.3), so there are simply more signal events in BSM scenarios. Importantly, the cross-section and Lorentz boost are the same for $\kappa_W = -1$, $\kappa_Z = 1$ and $\kappa_W = 1$, $\kappa_Z = -1$. The analysis was optimized for the latter case.

The VBS jets also provide a distinct kinematic signature (Fig. 4.4), namely two nearly back-to-back jets—i.e. a large absolute difference in pseudorapidity $|\Delta\eta_{jj}|$ —with a large combined invariant mass M_{jj} . In particular, the background processes fall off exponentially in M_{jj} whereas the signal process is more flatly distributed. Combined with the fact that the signal has a distinctly larger average value of $|\Delta\eta_{jj}|$ than background, these VBS characteristics form a strong handle for distinguishing signal from background.




Figure 4.2. Distribution of S_T plotted at the generator level for $\kappa_Z = +1$, demonstrating the Lorentz boost in the final state due to modifications to λ_{WZ} . Importantly, the distributions look identical for the case where $\kappa_W = +1$ and κ_Z is similarly varied.

4.3 The backgrounds

In general, background processes for this analysis need a fake¹ $H \rightarrow b\bar{b}$ merged jet, two fake¹ VBS jets, one lepton, and some E_T^{miss} . The leading background process for this analysis is $t\bar{t}$ production, wherein a top and antitop quark are produced and both decay to a bottom quark and W boson. One of the W bosons can decay to a real lepton and neutrino, and the fake VBS jets and $H \rightarrow b\bar{b}$ merged jet can come from some combination of the b quarks, quarks from one of the W bosons decaying hadronically, and possibly an extra gluon radiated by one of the incoming quarks (Fig. 4.5a). One of the incoming quarks can also radiate a W or Z boson (Fig. 4.5b), which presents additional opportunities for trickery. The largest sub-leading background comes from W+jets (Fig. 4.5c), where a W boson is produced along with some number of quarks or gluons, and single-top production (Fig. 4.5d), where only one top quark is produced. In W+jets, the W can give a real lepton and neutrino, while the additional jets cover

¹In principle, a background process could have a *real* Higgs boson with large p_T decay to $b\bar{b}$ or two *real* VBS jets, however SM processes with these kinds of signatures are so rare that they do not represent a significant background for this analysis.




Figure 4.3. The signal cross-section plotted as a function of enhancements to κ_W , keeping $\kappa_Z = +1$ (left) and to κ_Z , keeping $\kappa_W = +1$ (right). The black points are taken from MADGRAPH5_aMC@NLO, and the blue curve is the best fit of a 2nd order polynomial to those points.

the fake VBS and $H \rightarrow b\bar{b}$ signature. In single-top production, the top quark again decays to a W boson and b quark, so it can produce a signal-like final state much like $t\bar{t}$.

4.4 Event selection

4.4.1 Triggers and preselection

The single lepton in our signal final state gives a clean signature on which to trigger, so the aptly named “single lepton” triggers are applied and the “single lepton” datasets are used. These triggers are nearly 100% efficient for signal events because the lepton in the BSM final state is expected to have high p_T . A number of additional event-level filters are applied to remove detector noise and unphysical events [134].

Then, a loose selection, referred to as the “Preselection,” is applied to select the final state of interest. First, a loose selection is made on the combined VBS jet invariant mass: $M_{jj} > 500\text{GeV}$. Next, the PARTICLENET $X \rightarrow b\bar{b}$ score of the $H \rightarrow b\bar{b}$ fat jet candidate is required to be greater than 0.3. The event is also required to have no AK4 jets passing the




Figure 4.4. Event display for a simulated VBS WH event in the r - ϕ plane (left) and r - z plane (right), with the final-state particles (thick red lines) labeled. Two VBS jets can be seen very near to the beamline, going in either direction in z .

Medium DEEPJET working point. The event must furthermore have one and only one lepton with $p_T > 40\text{ GeV}$ that passes the Tight lepton ID. If there are any additional leptons that pass the Veto lepton ID, the event is vetoed—these leptons are not required to pass the p_T threshold. Finally, the event must have $S_T > 800\text{ GeV}$ (where S_T is defined in Eq. 4.3).

4.4.2 Signal region

The signal region for this analysis is defined on top of the Preselection with similar, but tighter selections. First, the S_T threshold is increased to $S_T > 900\text{ GeV}$. Then, the selections on the VBS jet variables are tightened to $M_{jj} > 600\text{ GeV}$ and $|\Delta\eta_{jj}| > 4$. Finally, the selections on the $H \rightarrow b\bar{b}$ fat jet candidate are made much more stringent, where the PARTICLENET $X \rightarrow b\bar{b}$ score is required to be greater than 0.9 and M_{SD} is required to be less than 150 GeV .

The background in this region is estimated using a data-driven technique, described in the next section, that the signal region was specifically designed to enable. Looser cuts are preferred in particular, as it was found that the variables used for the background estimation become correlated in a more restricted phase space. Therefore, the signal region selections are not optimized for maximal purity, though such a region can be formed.




Figure 4.5. Feynman diagrams for the leading and subleading backgrounds in this analysis. From top to bottom: $t\bar{t}$ production in the single-lepton final state with (a) an extra jet from a gluon or (b) an extra vector boson ($V = W$ or Z) radiated from one of the incoming partons; W +jets production in the single-lepton final state (c); and the production of a single top quark and W boson (d).

4.5 Background estimation

The background in the signal region is estimated using the “ABCD” method, where region A is the signal region, and regions B, C, and D have the $\Delta\eta_{jj}$ requirement, the M_{SD} requirement, and both inverted, respectively (Fig. 4.9). First, let the background yield in regions A, B, C, and D in Monte Carlo be defined as A_{MC} , B_{MC} , C_{MC} , and D_{MC} . Likewise, let the same yields in data be defined as A_{data} , B_{data} , C_{data} , and D_{data} . Under these definitions, the estimated background yield in region A, referred to as A_{data}^{pred} , can be computed with data as follows:

$$A_{data}^{pred} = B_{data} \times \frac{C_{data}}{D_{data}} \quad (4.4)$$




Figure 4.6. The VBS jet combined invariant mass M_{jj} (left) and absolute difference in pseudorapidity $|\Delta\eta_{jj}|$ (right) plotted after applying the Preselection.

where the same can be done in MC, yielding A_{MC}^{pred} . First, it can be seen in Fig. 4.11 that data and MC agree reasonably well in regions A, B, and C. In addition, it has been verified that the “transfer factor” that scales the actual yield in Region C to the estimated yield in region D is consistent within statistical uncertainty across data and MC:

$$\frac{C_{MC}}{D_{MC}} = 0.71 \pm 3.1\% \quad \frac{C_{\text{data}}}{D_{\text{data}}} = 0.71 \pm 11.0\%$$

The closure of the ABCD method described here is tested by comparing A_{MC}^{pred} to A_{MC} . This checks how closely the estimation in simulated events predicts the actual yield in simulation.

$$A_{MC}^{\text{pred}} = B_{MC} \times \frac{C_{MC}}{D_{MC}} = 129.7 \quad A_{MC} = 117$$

Comparing these two numbers, it is clear that the ABCD method for this analysis systematically over-predicts the background yield in the signal region. The difference between the predicted and actual yield in MC is therefore taken as the baseline systematic uncertainty of 11.3% on this method. However, the background composition is not consistent between the ABCD regions,




Figure 4.7. The variable $S_T = p_T(H) + p_T(\ell) + p_T^{\text{miss}}$ plotted after applying the Preselection.

and it is furthermore not known to high precision. The baseline systematic must be extended to cover the uncertainty in the background composition. In order to quantify this uncertainty, the yields for each of the next-to-leading backgrounds are varied within their respective statistical uncertainties, representing scenarios in which the contributions from each background are larger or smaller, and the closure of the method is recalculated. Based on how the closure varies for each scenario, a final systematic uncertainty of 13% is selected to cover the uncertainty in the background composition. Thus, the systematic and statistical uncertainties ϵ_{syst} and ϵ_{stat} are

$$\epsilon_{\text{syst}} \approx 13\%$$

$$\epsilon_{\text{stat}} = \frac{\sqrt{B_{\text{data}}}}{B_{\text{data}}} \oplus \frac{\sqrt{C_{\text{data}}}}{C_{\text{data}}} \oplus \frac{\sqrt{D_{\text{data}}}}{D_{\text{data}}} \approx 13.4\%$$

and the estimated background yield in the signal region is therefore given by

$$A_{\text{data}}^{\text{pred}} = 120 \pm 16.1 \pm 15.3$$

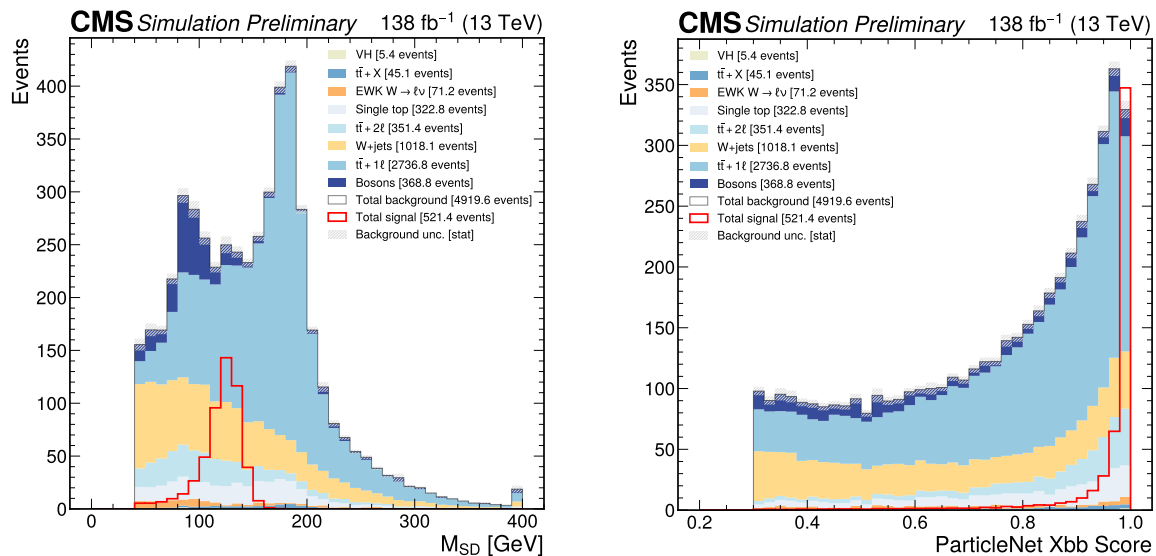


Figure 4.8. The $H \rightarrow b\bar{b}$ fat jet candidate soft drop mass (left) and PARTICLENET $X \rightarrow b\bar{b}$ score (right) plotted after applying the Preselection.

However, it can be seen in Fig. 4.10a that the ABCD variables $\Delta\eta_{jj}$ and M_{SD} are correlated. This correlation is not evident in MC closure because it happens to “cancel out,” leading to the mild systematic over-prediction quantified above. To improve the reliability and accuracy of the method, a correction is computed in MC and applied to the transfer factor $C_{\text{data}}/D_{\text{data}}$ as follows:

$$A_{\text{data}}^{\text{pred}} = B_{\text{data}} \times \frac{C_{\text{data}}}{D_{\text{data}}} \times \left(\frac{A_{\text{MC}}}{C_{\text{MC}}/D_{\text{MC}}} \right)$$

That is, the over-prediction of the ABCD method is measured in MC, i.e. the 13% computed before, but it is now used to adjust the transfer factor $C_{\text{data}}/D_{\text{data}}$ to account for the correlation between M_{SD} and $\Delta\eta_{jj}$. Moreover, in the $M_{SD} \geq 150\text{GeV}$ sideband, comparison between data and MC shows that the correlation is well-modeled by MC (Fig. 4.10b), thus validating this correction. This trivially makes the method close exactly in MC, but by applying it to ABCD in data, the predicted background yield in the signal region is made more accurate, as the correlation was previously biasing the result. Therefore, the final estimated background yield in the signal

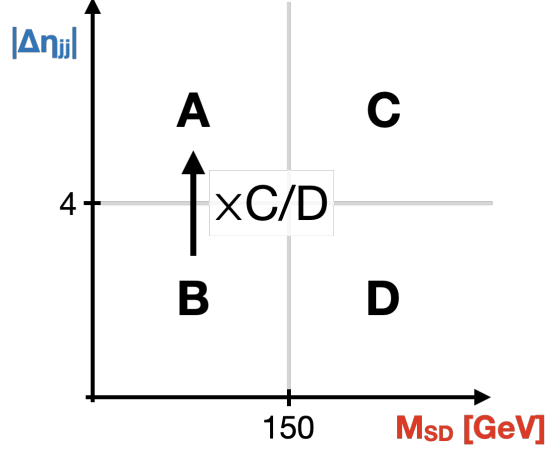


Figure 4.9. A sketch of regions A, B, C, and D used in the background estimation method.

region is given by

$$A_{\text{data}}^{\text{pred}} = 108 \pm 14.4 \pm 13.8$$

where the systematic and statistical uncertainties are kept as 13% and 13.4% respectively.

4.6 Results

The background yield in the signal region estimated from data and the signal yield predicted by Monte Carlo simulation are shown in Fig. 4.11 and tabulated in Table 4.1. Using these yields, and a thorough accounting of systematic uncertainties, we perform a maximum-likelihood fit [135] using the COMBINE statistical tool. Then, the exclusion significance and confidence level (CL) are extracted following the procedure described in Section 3.2 of Ref. [136]. In particular, we exclude the scenario where $\lambda_{WZ} = -1$ with an observed (expected) significance of 8.3σ (9.3σ), and we exclude all $\lambda_{WZ} < 0$ scenarios that were allowed by previous limits well beyond 5σ . These results are presented in Fig. 4.12a and 4.12b, respectively. This analysis, therefore, strongly indicates that κ_W and κ_Z have the same sign, representing another success of the Standard Model and the final determination of this property of the Higgs boson.

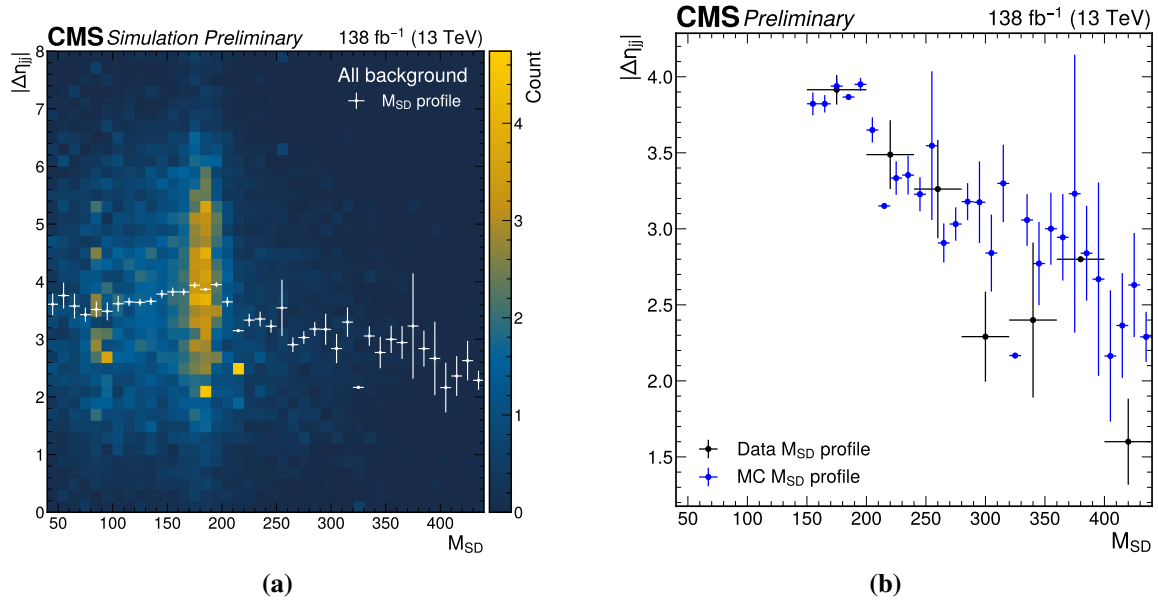


Figure 4.10. A 2-dimensional histogram of the ABCD variables M_{SD} and $\Delta\eta_{jj}$ with a profile of M_{SD} overlaid using only MC (left). The downward trend in the profile plot indicates that there is a minor correlation between M_{SD} and $\Delta\eta_{jj}$. The profile is compared between data and MC (right), and good agreement suggests that the correlation is at least well-modeled. This correlation can thus be addressed with a correction derived from MC when applying the method in data, such that the background yield in the signal region is more accurately predicted.

4.7 Acknowledgements

This chapter is a partial reproduction of the paper “Study of WH production through vector boson scattering and extraction of the relative sign of the W and Z couplings to the Higgs boson in proton-proton collisions at $\sqrt{s} = 13$ TeV,” submitted to PLB (arXiv:2405.16566). The work was made possible by the technical and administrative staff that operate and maintain the CERN accelerator complex, the LHC, CMS itself, and the worldwide LHC computing grid that provides data storage and processing capabilities for every LHC analysis.






Figure 4.11. The data and MC yields plotted in regions A, B, C, and D as a 1-dimensional histogram. In region A, the signal region, the background yield estimated from data is plotted instead of that predicted by MC. Overall, there is good agreement between data and MC in regions B, C, and D, with the difference between data and MC never exceeding 2σ .

Table 4.1. Table of the signal region yields from signal MC, background estimated from data, and observed data along with their respective statistical and systematic uncertainties. The systematic uncertainty for the signal quoted here is the sum in quadrature of all of the independent systematic uncertainties (percent errors) multiplied by the total yield. The observed yield is also tabulated.

Type	Yield	\pm stat.	\pm syst.
Signal	366	± 2.9	± 68.1
Background	108	± 14.4	± 13.8
Observed	130		



(a) 1D exclusion of $\lambda_{WZ} = -1$



(b) 2D exclusion of λ_{WZ} values

Figure 4.12. The negative-log-likelihood plotted as a function of the signal strength (μ) for $\lambda_{WZ} = -1$, with the 68% confidence interval labeled (a), and the exclusion significance plotted as a function of κ_W and κ_Z with the previous best CMS limits plotted as white capped error bars (b). It is clear from these two plots (a) the $\lambda_{WZ} = -1$ scenario is excluded beyond 5σ and (b) opposite-sign scenarios ($\lambda_{WZ} < 0$) within previous CMS limits are also excluded beyond 5σ .

Chapter 5

Measurement of κ_{VV} through VBS VVH production

5.1 Chasing couplings

The precise determination of the Higgs boson couplings to other particles is essential to establishing our understanding of the Higgs mechanism with empirical data. However, the more exotic the coupling, the more rare it is in nature. Among these are the HHVV and HHH couplings, whose modifiers are represented by κ_{VV} and κ_λ in the κ framework, respectively. At the time of writing, the value of κ_{VV} has been constrained to $[0.67, 1.38]$ and the value of κ_λ has been constrained to $[-1.24, 6.49]$ at the 95% CL [102]. Thus, compared to the κ_V coupling modifiers, for example, the values of κ_{VV} and κ_λ are relatively less well-measured, providing an exciting opportunity for advancing our understanding of the Higgs boson.

5.2 The signal

A less well-known process that includes κ_{VV} and κ_λ is the VBS production of HVV (Fig. 5.1). While rare, the cross-section of VBS HVV production grows almost quadratically with κ_{VV} (and more mildly for κ_λ), and the Higgs boson and vector boson receive a characteristic Lorentz boost for BSM values of either coupling modifier. This process has been proposed as a promising channel through which κ_{VV} and κ_λ could be measured [108], further reducing their

respective confidence intervals. By branching fraction, the VBS HVV all-hadronic final state is the most populous, although it has the largest background, followed by the single-lepton and dilepton final states. This chapter focuses on the all-hadronic final state.




Figure 5.1. Leading-order Feynman diagrams for VBS production of two vector bosons (V) and one Higgs boson, where the vector bosons decay hadronically and the Higgs boson decays specifically to b quarks. The HHV coupling κ_{VV} is denoted by a blue circle (●), and the HHH coupling κ_λ is denoted by a magenta circle (●).

5.2.1 Signal characteristics

The VBS quarks provide the usual signature: two nearly back-to-back VBS jets have large $|\Delta\eta_{jj}|$ and M_{jj} (Fig. 5.2). The boost from BSM values of κ_{VV} is also evident when the VVH system is reconstructed as three AK8 jets (Fig. 5.3). Moreover, a graph neural network called PARTICLENET can be used to classify each AK8 jet, with negligible misidentification for signal, providing a strong handle for removing the dominant QCD background. The PARTICLENET regressed mass M_{PNet} is used to estimate the mass of each AK8 jet, as it has been shown to have better resolution than the other methods, and the clear resonance peaks around the Higgs boson and vector boson masses can be used to separate signal from background, which, in general, does not have such resonances (Fig. 5.4).

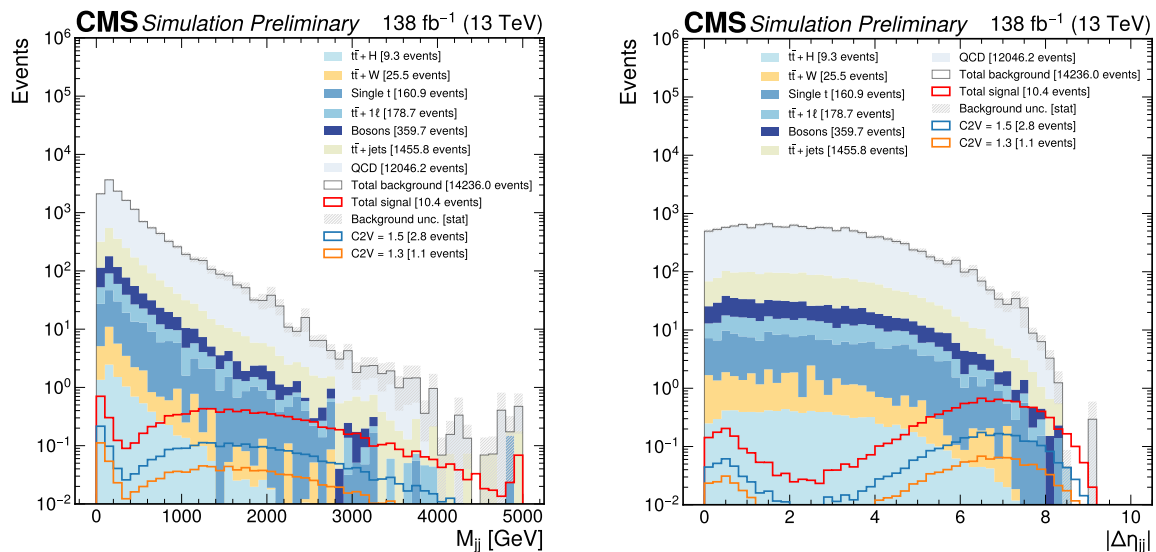


Figure 5.2. The invariant mass M_{jj} (left) and pseudorapidity separation $\Delta\eta_{jj}$ (right) of the VBS jets plotted after the Preselection is applied. The main signal MC ($\kappa_{VV} = 2$) is plotted in red alongside $\kappa_{VV} = 1.5$ and $\kappa_{VV} = 1.3$ for comparison.

5.3 The backgrounds

The main background in this analysis comes from QCD multijet production (Fig. 5.5a), wherein the quarks and gluons in the colliding protons interact and produce more quarks and gluons. This is the most common process at the LHC, and it dominates the background for many analyses that target an all-hadronic final state. The largest sub-leading backgrounds are $t\bar{t}$ production (Fig. 5.5b) and multi-boson production in all-hadronic final states. All of these backgrounds simply produce many jets that can, by random chance, pass our $H \rightarrow b\bar{b}$ selections, or have genuine vector bosons that decay hadronically.

5.3.1 QCD PARTICLENET resampling

From Fig. 5.6, it is clear that the data and MC distributions for the PARTICLENET scores used in this analysis do not fully agree. This is not an issue for the final result, as the background estimation is fully data-driven. However, for training any classifier for this analysis, such as a boosted decision tree (BDT) or deep neural network (DNN), better modeling would result in




Figure 5.3. The p_T of the $H \rightarrow b\bar{b}$ (left), leading $V \rightarrow qq$ (center), and trailing $V \rightarrow qq$ (right) AK8 jets plotted after the Preselection is applied. The main signal MC ($\kappa_{VV} = 2$) is plotted in red alongside $\kappa_{VV} = 1.5$ and $\kappa_{VV} = 1.3$ for comparison.




Figure 5.4. The PARTICLENET regressed masses M_{PNet} of the $H \rightarrow b\bar{b}$ (left), leading $V \rightarrow qq$ (center), and trailing $V \rightarrow qq$ (right) AK8 jets plotted after the Preselection is applied. The main signal MC ($\kappa_{VV} = 2$) is plotted in red alongside $\kappa_{VV} = 1.5$ and $\kappa_{VV} = 1.3$ for comparison.

more events passing the Preselection, and therefore being available for training.

To address this, the following resampling procedure is used to correct the PARTICLENET score shape for QCD MC. Only the scores for the QCD jets undergo this procedure, and the rest of the minor backgrounds are subtracted from the data distributions used for the resampling. First, the following assumptions are made:

1. A mass-decorrelated PARTICLENET score for a given AK8 jet should be fundamentally described by a probability density function \mathcal{P} , namely $\mathcal{P}_{X \rightarrow b\bar{b}}$ and $\mathcal{P}_{X_W \rightarrow qq}$ for the




Figure 5.5. Feynman diagram for QCD multijet production (left) and $t\bar{t}$ production in the all-hadronic final state (right), the leading and subleading backgrounds, respectively, for the VBS VVH all-hadronic analysis.




Figure 5.6. The PARTICLENET scores for the $H \rightarrow b\bar{b}$ (left), leading $V \rightarrow qq$ (center), and trailing $V \rightarrow qq$ (right) AK8 jets plotted after the analysis objects are selected. The main signal MC ($\kappa_{VV} = 2$) is plotted in red alongside $\kappa_{VV} = 1.5$ and $\kappa_{VV} = 1.3$ for comparison.

PARTICLENET $X \rightarrow b\bar{b}$ and $X_W \rightarrow qq$ scores, respectively.

2. \mathcal{P} should be the same regardless of how many AK8 jets are in the event.
3. The PARTICLENET score is not correlated to the jet kinematic properties used later in the analysis.

Assumption 1 is the common statistical interpretation of a histogram; that is, \mathcal{P} can be approximated by histogramming the PARTICLENET score for every AK8 jet in an event and normalizing it to unity. Assumption 2 is less trivial. In our case, we need to verify that events with 3 AK8 jets to contain the same kinds of AK8 jets as an event with 2 AK8 jets. This is not guaranteed, as

events with different AK8 jet multiplicity may involve different physics. For QCD events in this analysis, however, it can be seen in Fig. 5.7 that Assumption 2 holds, as the *PDFs* in 3 AK8 jet events approximately match the 2 AK8 jet events in data (with non-QCD events taken from MC subtracted out). Therefore, the properly binned \mathcal{P} derived in data for 2-AK8-jet events can be sampled to replace the PARTICLENET scores in 3-AK8-jet events. In order to account for the correlation in this resampling, we bin the PARTICLENET $X \rightarrow b\bar{b}$ score in p_T .

The PARTICLENET $X_W \rightarrow qq$ score, on the other hand, must be binned in the $X \rightarrow b\bar{b}$ score and p_T , since the $X \rightarrow b\bar{b}$ score is included indirectly in the calculation of the $X_W \rightarrow qq$ score. The leading $X \rightarrow b\bar{b}$ score is also picked first, thus potentially further biasing the correlation among the scores. The replacement is performed for each AK8 jet via the following prescription:

1. Get AK8 jet p_T
2. Sample $\mathcal{P}_{X \rightarrow b\bar{b}}$ for the appropriate p_T bin
3. Replace AK8 jet $X \rightarrow b\bar{b}$ score with one randomly sampled from $\mathcal{P}_{X \rightarrow b\bar{b}}$
4. Get $\mathcal{P}_{X_W \rightarrow qq}$ for the appropriate p_T , $X \rightarrow b\bar{b}$ score bin
5. Replace AK8 jet $X_W \rightarrow qq$ score with one randomly sampled from $\mathcal{P}_{X_W \rightarrow qq}$

After replacing the AK8 jet scores in events with 3 AK8 jets, the PARTICLENET score distribution agreement between data and MC improves significantly (Fig. 5.8), doubling the number of raw QCD MC events that pass the Preselection.

5.4 Event selection

5.4.1 Triggers and preselection

Because the entire final state is composed of jets, the jet H_T triggers are applied, where H_T is “hadronic transverse energy,” or the sum of the p_T of every jet in the event. Moreover,

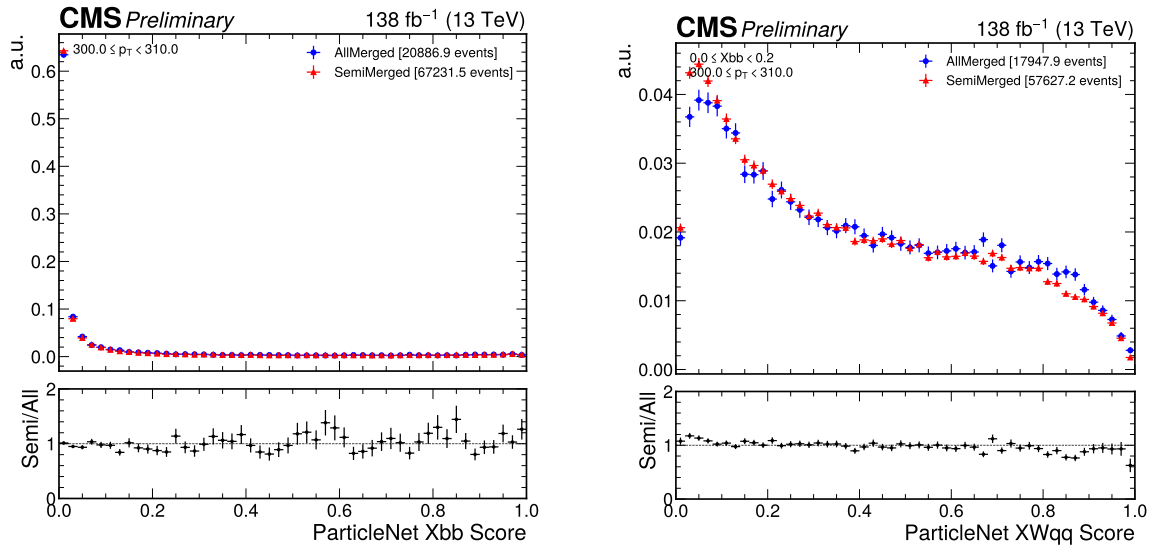


Figure 5.7. The approximation of the PARTICLENET probability density functions plotted for the $X \rightarrow b\bar{b}$ score in one bin of p_T (left) and the $X_W \rightarrow qq$ score in one bin of p_T and $X \rightarrow b\bar{b}$ (right).

since the jets considered in this analysis are expected to be highly boosted, these triggers are 100% efficient for signal events. The same event-level filters used to remove detector noise and unphysical events in the analysis described in the previous chapter are applied also in this analysis.

Also like the previous analysis, a loose selection referred to as the “Preselection” is applied to select the final state of interest. First, events are required to have zero leptons passing the veto lepton ID. This orthogonalizes the all-hadronic channel from the other channels that are analyzed separately. The event must have at least three AK8 jets, where the leading AK8 jet is required to have $p_T > 500\text{ GeV}$ (while for all of them the p_T threshold is 250 GeV) such that the analysis is on the H_T -binned trigger threshold. Finally, the event is also required to have two AK4 jets, following the VBS jet selection scheme detailed in Section 3.4.1. Then, the AK8 jet with the highest PARTICLENET X_{bb} score is selected as the $H \rightarrow b\bar{b}$ candidate. The next two AK8 jets are selected as the $V \rightarrow qq$ candidates, sorted (leading vs. trailing) by p_T . With the AK8 jets thus tagged, a very loose selection is applied to the PARTICLENET scores: the $H \rightarrow b\bar{b}$ candidate is required to have a PARTICLENET $X \rightarrow b\bar{b}$ score greater than 0.5, while the $V \rightarrow qq$




Figure 5.8. The PARTICLENET scores for the $H \rightarrow b\bar{b}$ (left), leading $V \rightarrow qq$ (center), and trailing $V \rightarrow qq$ (right) AK8 jets plotted after the analysis objects are selected and the QCD resampling is applied. The main signal MC ($\kappa_{VV} = 2$) is plotted in red alongside $\kappa_{VV} = 1.5$ and $\kappa_{VV} = 1.3$ for comparison.

candidates are each required to have a PARTICLENET $X_{W \rightarrow qq}$ score greater than 0.3.

5.4.2 ABCDNET

The signal region would preferably be accompanied by a valid data-driven background estimation method while also being optimized for maximal significance, since the dominant background is multijet QCD. Because this is a common analysis scenario, a machine learning (ML) method has been proposed to satisfy both requirements at once [137]. This method has been used already in some published work, e.g. Ref. [138]. In particular, a DNN is trained to serve as one of the “arms” of a traditional ABCD background estimation. Critically, a $d\text{Corr}^2$ term is added to the loss function (\mathcal{L}) that trains the DNN to be decorrelated with the other arm:

$$\mathcal{L}[f(\vec{x})] = \mathcal{L}_{\text{BCE}}[f(\vec{x}, y)] + \lambda d\text{Corr}_{y=0}^2[f(\vec{x}), X_0] \quad (5.1)$$

where \vec{x} is the input vector, y is the truth label (1 for signal, 0 for background), λ is a tunable parameter for controlling the size of the decorrelation term, and X_0 is the decorrelation target, which the analyzer is free to choose. Binary Cross Entropy (\mathcal{L}_{BCE}) is used in our analysis, however any loss function \mathcal{L} could in principle be used in its place. The $d\text{Corr}^2$ term is the

“distance correlation,” a statistical quantity that measures the dependence of two variables f and g , based on the “distance covariance” dCov^2 between them:

$$\text{dCov}^2[f, g] = \langle |f - f'| \times |g - g'| \rangle + \langle |f - f'| \rangle \times \langle |g - g'| \rangle - 2\langle |f - f'| \times |g - g''| \rangle \quad (5.2)$$

$$\text{dCorr}^2[f, g] = \frac{\text{dCov}^2[f, g]}{\text{dCov}[f, f]\text{dCov}[g, g]} \quad (5.3)$$

For this analysis, a DNN hereafter referred to as ABCDNET was trained to be decorrelated from $|\Delta\eta_{jj}|$. While there are many configurations of this technique that were considered, we found that using $|\Delta\eta_{jj}|$ as the decorrelation target provided the most stable result and the best closure for this analysis. A simple architecture was selected for ABCDNET: 3 hidden layers with 64 nodes each. Additional hyperparameters are tabulated in Table 5.1. A total of 13 features (Fig. 5.9) are provided to ABCDNET as inputs:

- $H \rightarrow b\bar{b}$ candidate $p_T, \eta, \phi, M_{\text{PNet}}$
- Leading $V \rightarrow qq$ candidate $p_T, \eta, \phi, M_{\text{PNet}}$
- Trailing $V \rightarrow qq$ candidate $p_T, \eta, \phi, M_{\text{PNet}}$
- M_{jj}

In short, the 4-vectors for the Higgs candidate and two V candidates are supplied, along with M_{jj} . While the decorrelation term in the loss is able to train the DNN to decorrelate even fundamentally correlated variables, like M_{jj} and $\Delta\eta_{jj}$, adding any additional VBS variables harmed the final closure. In addition, the background and signal event weights were renormalized such that the total integral for each were equal to the total number of raw background events. The input features were all normalized to be of order 1—this helps the model converge. Specifically, the p_T is log-normalized ($p_T \rightarrow \log(p_T)$) and the other variables are normalized as follows:

$$x \rightarrow \frac{x - x_{\min}}{x_{\max} - x_{\min}} \quad (5.4)$$

That is, the range of a variable x is normalized such that a chosen minimum (x_{\min}) and maximum (x_{\max}) value are scaled to 0 and 1 respectively. This range is selected to be within a reasonable window; for example, M_{jj} is normalized such that 0 remains 0 and 3000 is scaled to 1.

Table 5.1. The hyperparameters for ABCDNET. Many values of λ were tried, where the best value was determined by comparing the correlation between the ABCDNET discriminator and $\Delta\eta_{jj}$, as well as comparing the ABCD closure directly.

Parameter	Setting
Number of hidden layers	3
Hidden layer size	64
Activation function	Leaky ReLU
Learning rate (constant)	0.001
Test/train split	80/20
Number of training batches	10
Number of testing batches	5
DisCo λ	30

ABCDNET was trained for 3000 epochs with a constant learning rate. Once training was complete, the model was selected before the average testing and training loss start to diverge. Notably, the decorrelation term tends to overfit faster than BCE, and preference must be given to less overfitting in the decorrelation than in the overall performance. For this analysis, epoch 700 was deemed optimal (Fig. 5.10a). The selected model shows no signs of overfitting and sufficient performance (Fig. 5.10b). Moreover, it can be seen in Fig. 5.11 that the ABCDNET discriminant (Fig. 5.12) and $|\Delta\eta_{jj}|$ are indeed decorrelated.

5.4.3 Signal region

Ultimately, the signal region for this analysis is formed from cuts on five variables: ABCDNET score, $|\Delta\eta_{jj}|$, and PARTICLENET scores for the three AK8 jets. A brute-force scan was performed, where over fifty thousands regions were tested. The regions were ranked by a rough estimation of the significance: S/\sqrt{B} , where S is the signal yield taken from MC and B is the background yield predicted using data via the ABCD method described in the next section. The final signal region was required to have at least 0.5 predicted background events, as cutting

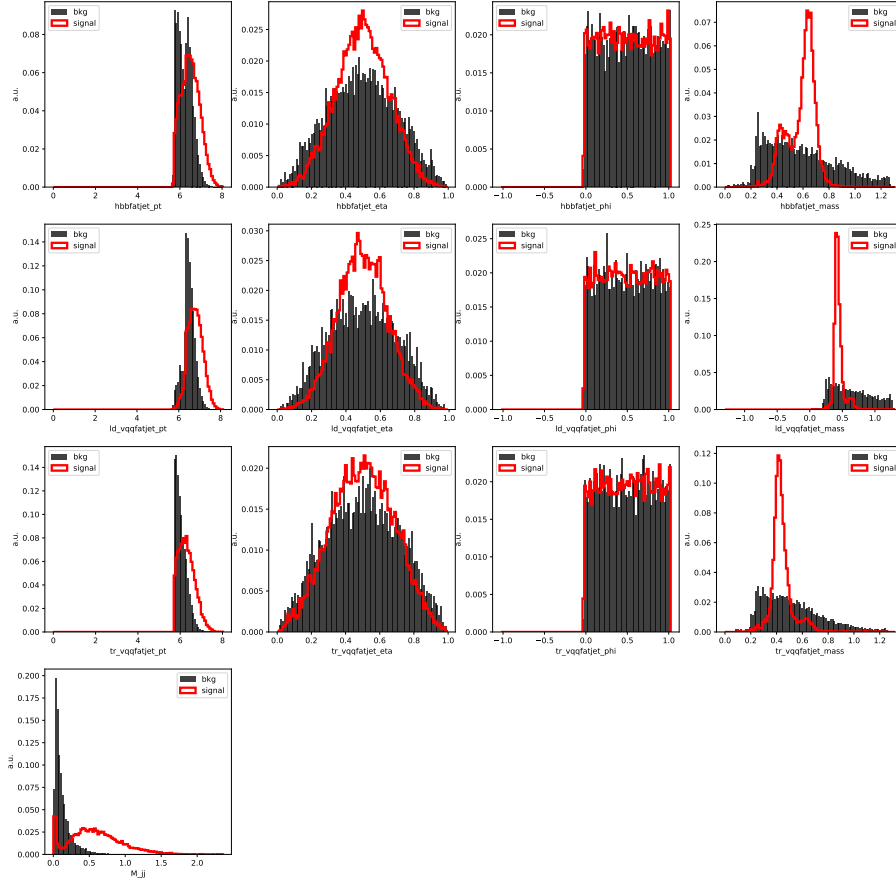


Figure 5.9. The input features for ABCDNET plotted as histograms normalized to unity. It can be seen that the features have been re-scaled to values of order 1.

tighter on the signal region variables resulted in larger uncertainties. Ultimately, the following signal region was determined:

- $\text{ABCDNET} > 0.89$
- $|\Delta\eta_{jj}| > 5$
- $X_{bb}(H \rightarrow b\bar{b}) > 0.8$
- $X_{Wqq}(\text{leading } V \rightarrow qq) > 0.8$
- $X_{Wqq}(\text{trailing } V \rightarrow qq) > 0.7$

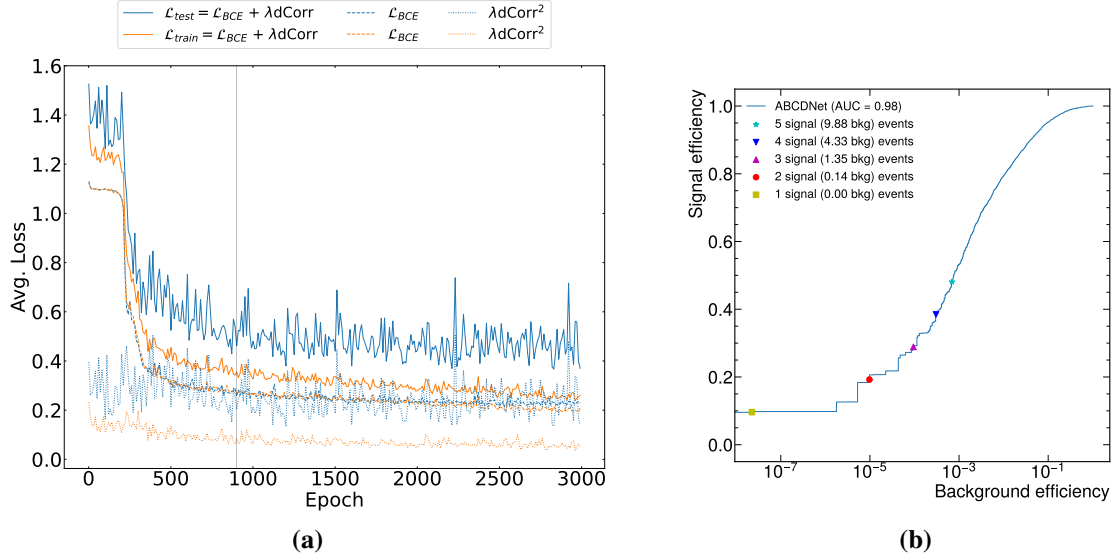


Figure 5.10. The average loss at each epoch (left) and ROC curve for the selected model (right). The total loss, which is the sum of the Binary Cross Entropy (BCE) plus the decorrelation term, is plotted as a solid line, the BCE is plotted as a dashed line, and the decorrelation term is plotted as a dotted line. The testing loss is plotted in blue and the training loss is plotted in orange.

5.5 Background estimation

The background in the signal region is estimated entirely from data by using the ‘‘ABCD’’ method, where regions A, B, C, and D are illustrated in Fig. 5.13. These regions are formed by selecting two cuts in the signal region (region A), then inverting them to form regions B, C, and D. So long as the two cuts are independent, the following method holds. As in the previous analysis, the background yield in regions A, B, C, and D in Monte Carlo are defined as A_{MC} , B_{MC} , C_{MC} , and D_{MC} . Likewise, let the same yields in data be defined as A_{data} , B_{data} , C_{data} , and D_{data} . Again, the estimated background yield in the signal region A_{data}^{pred} can be computed using Eq. 4.4, reproduced here for convenience:

$$A_{data}^{pred} = B_{data} \times \frac{C_{data}}{D_{data}}$$

where the same can be done in MC, yielding A_{MC}^{pred} . The two cuts used to define the ABCD method are $ABCDNET > 0.89$ and $|\Delta\eta_{jj}| > 5$. Because ABCDNET was trained to be decorrelated with

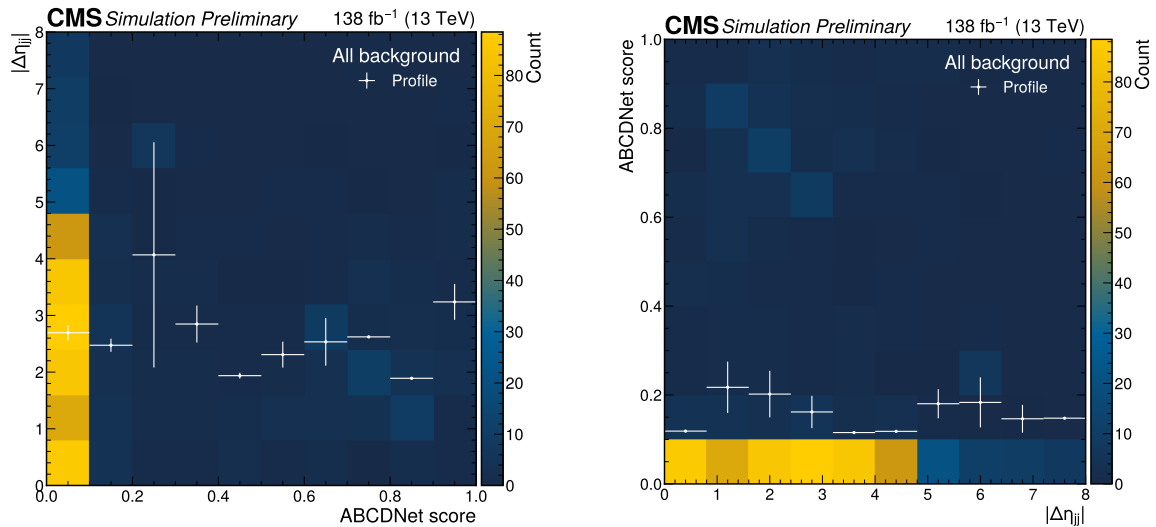


Figure 5.11. A two-dimensional histogram binned in ABCDNET score and $\Delta\eta_{jj}$. The one-dimensional profile of the x-axis is overlaid such that the correlation between the two variables appears as a trend in the profile. The signal region selections on the AK8 jet PARTICLENET scores are applied, so the left plot is equivalent to Fig. 5.13.

$\Delta\eta_{jj}$, the method holds by construction. Now, it can be seen in Fig. 5.15 that data and MC agree reasonably well in regions A, B, and C. The statistical uncertainty ϵ_{stat} on the method is simply the propagation of statistical uncertainties on the data yields for regions B, C, and D:

$$\epsilon_{\text{stat}} = \frac{\sqrt{B_{\text{data}}}}{B_{\text{data}}} \oplus \frac{\sqrt{C_{\text{data}}}}{C_{\text{data}}} \oplus \frac{\sqrt{D_{\text{data}}}}{D_{\text{data}}} \approx 34\% \quad (5.5)$$

Table 5.2. Data yields and region A prediction for the control region used for the ABCD closure test. The region A yield is kept blind, while A_{pred} is reported.

Region	Data Yield	Prediction
A	–	1.07
B	10	
C	72	
D	672	

We can immediately assess the closure of the method using only simulated events by

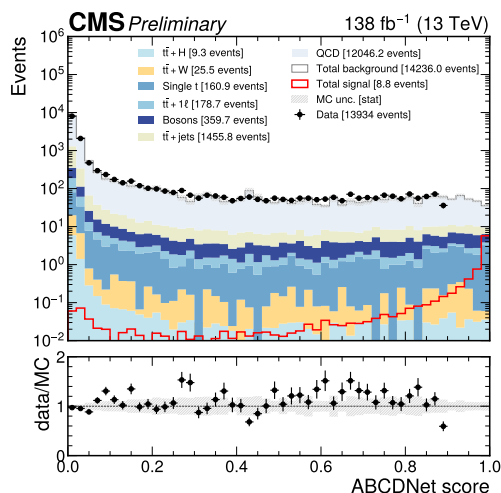


Figure 5.12. The ABCDNET score plotted in data and MC at the Preselection. The data beyond the cut used in the signal region is kept blind.

comparing A_{MC}^{pred} and A_{MC} :

$$A_{MC}^{\text{pred}} = B_{MC} \times \frac{C_{MC}}{D_{MC}} = 1.12 \pm 0.29 \quad A_{MC} = 2.34 \pm 0.82$$

However, the conclusion from this test is not clear, given the large uncertainty on the MC prediction in region A. Instead, a systematic uncertainty on the method ϵ_{syst} can be derived using MC as follows:

$$\epsilon_{\text{syst}} = \frac{2}{A_{MC} + A_{MC}^{\text{pred}}} \sqrt{\left(\frac{A_{MC}^{\text{err}}}{A_{MC}}\right)^2 + \left(\frac{\epsilon_{\text{stat}}^{\text{MC}}}{A_{MC}^{\text{pred}}}\right)^2} \approx 25\% \quad (5.6)$$

This effectively takes the relative error on the non-closure in MC as the systematic uncertainty, where A_{MC}^{err} is the statistical error on the MC yield and $\epsilon_{\text{stat}}^{\text{MC}}$ is the equivalent of Eq. 5.5 for MC.

Ultimately, the ABCD method will be applied exclusively in data, so showing good closure only in MC would not fully ensure the validity of the method in data. Moreover, the MC statistics are limited in the signal region, so the exact quality of the closure is not clear from MC alone. In order to verify that the method is valid in data, the sidebands B, C, and D can be divided into sub-regions as shown in Fig. 5.14b. The ABCD method should still hold in these

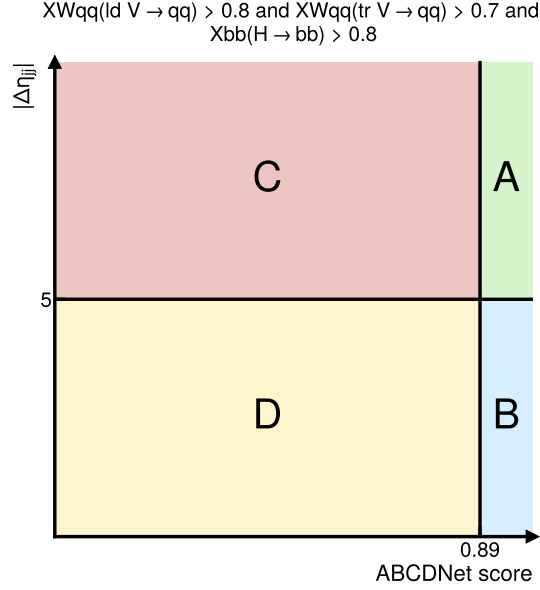


Figure 5.13. A cartoon of the ABCD regions. The signal region selections on the AK8 jet PARTICLENET scores are applied to all regions. Region B has the $\Delta\eta_{jj}$ cut inverted, region C has the ABCDNET score inverted, and region D has both cuts inverted.

sub-regions, and because regions B, C, and D exclude the signal region, closure can be evaluated in data. In total, 5 such checks are done, and the results are tabulated in Table 5.3. In each test, good closure is seen in data, supporting the validity of the method.

5.6 Results

We fit the signal and predicted background yield in the signal region to data for the analysis described here as well as all other channels. The B, C and D region data yields are included in the fit together with their respective signal contamination. We set upper limits at the 95% CL on the cross-section of the VBS VVH process for each κ_{VV} generated data point. The upper limits are calculated using the AymptoticLimits method in the COMBINE toolkit [128]. The full list of systematic uncertainties are included as nuisance parameters during the fitting procedure. The result of the fit is shown in Fig. 5.16a and 5.16b for the all-hadronic channel and the combination of all channels, respectively. The regions where the expected limit is smaller than the theoretical prediction on the cross-section are taken to be excluded values for κ_{VV} .

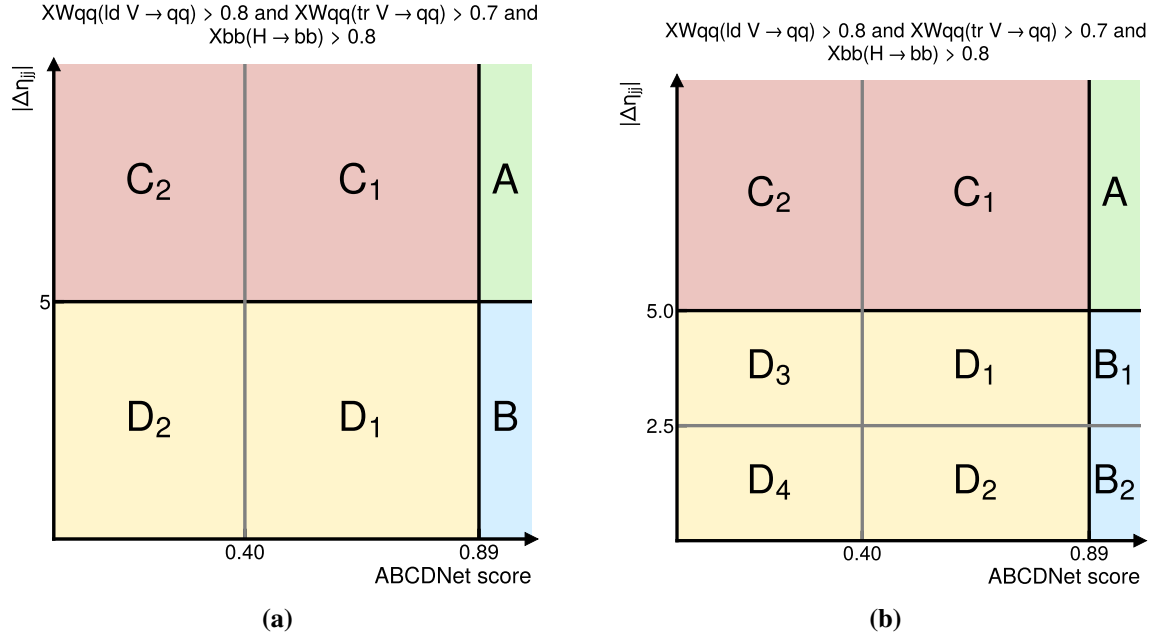


Figure 5.14. The ABCD regions used for closure tests. Again, the cuts on the AK8 jet PARTICLENET scores are applied to all regions. Regions C and D can each be split into two sub-regions by introducing a looser cut on the ABCDNET score (a). Region B can similarly be split into two sub-regions by introducing a looser cut on $\Delta\eta_{jj}$, which also divides D into 4 sub-regions total (b).

Specifically, the all-hadronic channel places a limit on the HHVV coupling to be within -0.03 and 2.04 times the SM, and the combined limit further limits it to within 0.23 and 1.78 times the SM.

5.7 Next steps

This work is still going through the rigorous CMS Collaboration approval process, wherein the methods and conclusions presented in this chapter are thoroughly checked for validity and accuracy. All of the results presented here use only simulated events—or, for the background estimation, recorded data events that do not enter the signal region. Soon, the analysis will be “unblinded,” meaning the yield of recorded data in the signal region will be seen for the first time. Already, the combination of all VBS VVH channels are expected to give a limit that is competitive with the current best result. In Run 3, and beyond, the combination of




Figure 5.15. Regions A and C (left) and regions B and D (right) used in the ABCD method.

VBS VVH with the results from other analyses will yield a precise measurement of the strength of the HHVV coupling.

5.8 Acknowledgements

This chapter is a partial reproduction of a paper being prepared for submission for publication. The work was made possible by the technical and administrative staff that operate and maintain the CERN accelerator complex, the LHC, CMS itself, and the worldwide LHC computing grid that provides data storage and processing capabilities for every LHC analysis.

Table 5.3. Data yields and signal-like region prediction for the ABCD closure tests. Regions B_i , C_i , and D_i are described in Fig. 5.14.

Region	Data Yield	Prediction	Region	Data Yield	Prediction
B_1	3 ± 1.7	5.6 ± 2.5	D_1	27 ± 5.2	24 ± 4.6
B_2	7 ± 2.7		D_2	34 ± 5.8	
D_1	27 ± 5.2		D_3	255 ± 16	
D_2	34 ± 5.8		D_4	356 ± 19	

Region	Data Yield	Prediction
C_1	5 ± 2.2	7.1 ± 1.7
D_1	27 ± 5.2	
C_2	67 ± 8.2	
D_3	255 ± 16	

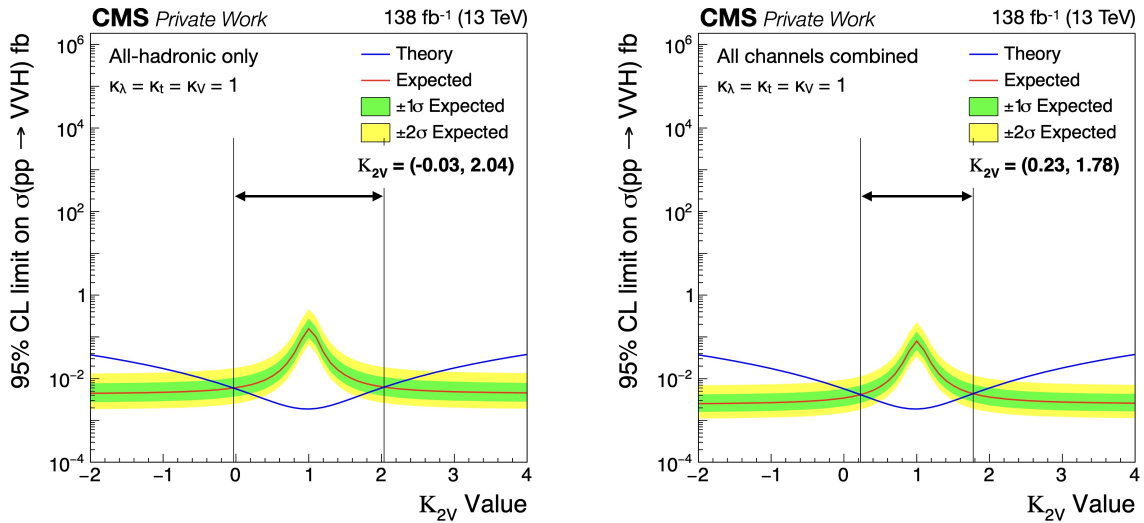


Figure 5.16. The 95% confidence level limit plotted as a function of κ_{VV} for the all-hadronic channel (left) and all channels combined (right).

Chapter 6

Line Segment Tracking

6.1 The HL-LHC computing challenge

With the massive increase in pileup in the HL-LHC era, each recorded event is massively more complex, making all steps of data processing much more expensive. Current projects show that, in fact, computing demands will exceed the resources that CMS is able to provide, even with annual increases (Fig. 6.1a). We must therefore make use of novel hardware or develop more efficient algorithms to make HL-LHC operations at all possible at CMS. In particular, reconstruction will represent 61% of CPU usage at CMS when the HL-LHC turns on (Fig. 6.1b), and a large fraction of that compute is dedicated to track reconstruction.




Figure 6.1. The HL-LHC CPU usage projections at CMS as a function of time (left) and broken down by computing tasks (right), from Ref. [139].

6.1.1 Track reconstruction

Because many thousands of particles are produced simultaneously in each bunch crossing, individual tracks need to be recognized out of dense clouds of x-y-z points called “hits” (Fig. 6.2). This is made even more challenging with the HL-LHC, where tens of pileup collisions becomes hundreds. Nevertheless, individual particle tracks *must* be reconstructed because they contain critical information about what was produced in the collision. This is accomplished in two steps: track finding and track fitting. First, a track-finding algorithm identifies each set of hits that were likely to have been generated by the same particle—these are called track “candidates.” Then, a track-fitting algorithm takes each track candidate and fits a trajectory to it, from which it can determine key properties of the particle like its charge and momentum.




Figure 6.2. Simple illustration of particles (gray lines) originating from the primary vertex (red star) and a secondary vertex (blue star), each leaving hits (black dots) in a hypothetical multi-layered tracker.

Better track-fitting yields more accurate “track parameters” that are crucial for downstream analysis. For example, those same algorithms that use the presence of displaced tracks to identify longer-lived particles rely on the accuracy of the input features. While not a central topic of this chapter, there is already a well-optimized solution for track-fitting [140].

Better track-finding yields a higher quantity and quality of reconstructed tracks, enabling

more diverse and precise analysis. Moreover, many interesting physics processes involve particles that decay in flight after traveling millimeters to centimeters (and beyond in some cases), resulting in “displaced” tracks that do not point to a proton-proton collision. The Higgs boson, a primary object of interest, can only be detected by identifying its decay products, some of which have longer lifetimes [114]. Many new physics candidates are also expected to have long lifetimes (e.g. Ref. [141]), so it is imperative that track-finding algorithms are robust against edge cases—in addition to displaced production vertices, tracks can also have holes or significant overlap. Finally, with the massive increase in pileup, track finding also becomes a problem of computational scalability, adding yet another dimension to the problem of track finding: proposed algorithms must meet or exceed critical efficiency milestones while being robust to vital edge cases and delivering massive gains in throughput all at once.

Traditional track-finding algorithms, namely the Kalman filter, proceed sequentially, building tracks from the innermost to the outermost layer of the silicon tracker. Moreover, this means that track-finding at CMS is currently completely reliant on the pixel detector, which is more prone to failures due to its proximity to the beamline, and displaced tracks that begin in the outer tracker are more easily missed.

6.2 The line segment tracking algorithm

Proposed originally as a drastic redesign of the silicon tracker layers, the line segment tracking (LST) algorithm makes use of the bi-layer “ p_T -modules” that will replace the single-layer tracking modules currently used in the outer tracker. Each of these modules will have two silicon sensors spaced a few millimeters apart [97], allowing for a rough estimation of a throughgoing particle’s p_T (Fig. 6.3). That is, particles with high- p_T will have two hits within a small window, so particles with two hits outside of that window have low- p_T and can most likely be rejected. This massively reduces the occupancy (number of hits) in the tracker due to pileup, which produces mostly low- p_T particles.

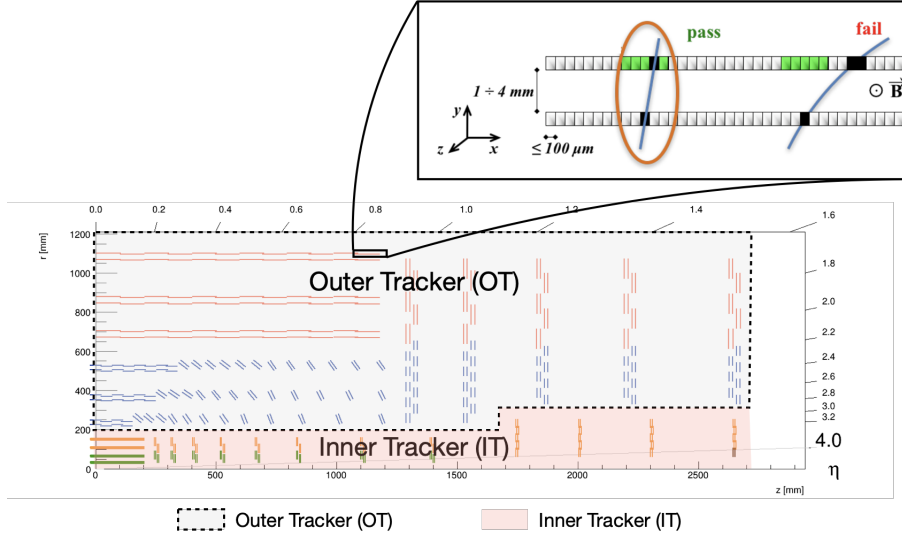


Figure 6.3. An illustration of the CMS Phase-II silicon tracker with a diagram of a p_T -module from Ref. [142].

Rather than build every track hit-by-hit, LST builds track segments of increasing size in parallel, leveraging GPUs for a massive speedup. Starting in the outer tracker, LST first builds “mini-doublets” (MDs) in each p_T -module, filtering out low- p_T tracks. Then, line segments (LSs) are built from each pair of MDs in adjacent tracker layers, and each LS is required to contain MDs with consistent p_T estimates. Progressively longer track segments are built in this way: “triplets” (T3s) are formed from LSs that share an MD, and “quintuplets” (T5s) are formed from T3s that share an MD. At each step, the track segments are required to be approximately consistent with a single helical trajectory. Finally, the track segments in the pixel layer (pLSs) are taken from an upstream algorithm and matched to the LST track segments, resulting in four mutually exclusive collections of track candidates (TCs). First, pLSs are matched to T5s, forming pT5s. Next, the unmatched T5s are set aside, and the T3s that are not in a pT5 are matched to pLSs, forming pT3s. Last, all unused pLSs are collected and considered as track candidates, in order to recover tracks from particles outside of the outer tracker acceptance region. A complete taxonomy of the LST objects is presented in Table 6.1.

We benchmark LST with a variety of MC simulation samples. Primarily, we use simulated $t\bar{t}$ events, which have displaced tracks (from b quarks) and multiple jets, and therefore a high

Table 6.1. The steps of the LST algorithm are shown in order of execution, starting with track-segment building (Steps 1 to 4) followed by track-candidate selection (Steps 5 to 8). Step 0 is performed by a preceding iteration of the CMS track-finding algorithm. Each LST step is implemented as a separate kernel, where the track objects of interest are built in parallel.

Step	Track segment	Description
0	Pixel seeds (pLSs)	Track segments from the inner tracker
1	Mini-doublets (MDs)	Two hits in a p_T -module
2	Line segments (LSs)	Two MDs in nearby modules
3	Triplets (T3s)	Two LSs that share a common MD
4	Quintuplets (T5s)	Two T3s that share a common MD
Step	Track candidate	Description
5	pT5s	T5s matched to a pixel seed
6	T5s	T5s that are not matched to a pixel seed
7	pT3s	T3s that are not in a pT5, but are matched to a pixel seed
8	pLS	Pixel seeds that are not already in a pT5 or pT3

tracker occupancy, with a pileup of 200 (PU200) and the Phase-II tracker geometry. In order to measure the performance of LST, we look at each simulated track and try to match it to a TC: if more than 75% of the hits in a TC belong to a single simulated track, it is considered as real. Then, the efficiency is defined as the fraction of simulated tracks that are matched in this way to a TC, whereas the fake rate is defined as the fraction of TCs that are not matched to any simulated tracks. These two metrics are plotted in Fig. 6.4, where it is clear that the fake TCs in the barrel are mostly T5s, which also contribute the largest fraction of the overall efficiency. In the baseline version of LST, there is therefore an opportunity for massive improvement.

6.3 Improving LST with machine learning

6.3.1 Training

We trained a DNN on the set of T5s built by LST running on 175 $t\bar{t}$ PU200 events with the selections on a custom heuristic on the quality of a circular fit to the T5 in the r - ϕ plane removed based on the assumption that the DNN could better use the same information. In total, this yielded approximately 840 000 real T5s and 1.26 million fake T5s shuffled into two datasets,

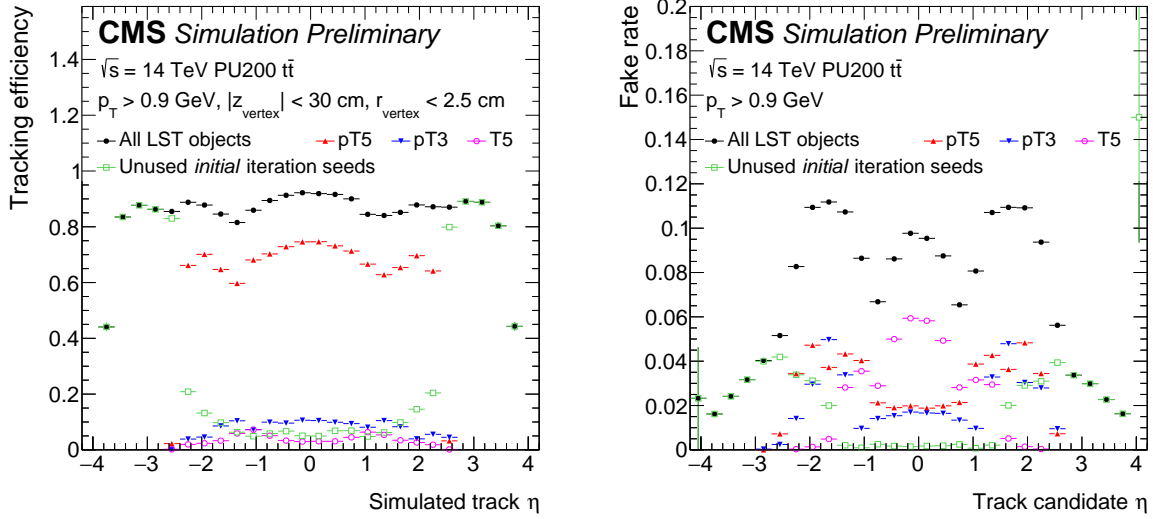


Figure 6.4. The LST track-finding efficiency (left) and fake rate (right) plotted as a function of the pseudorapidity η of the simulated track and track candidate, respectively. For both plots, five histograms are overlaid: all track candidates (black), pT5s (red), pT3s (blue), T5s (magenta), and pLSs that are not used in a pT5 or pT3 (green).

80% for training and 20% for testing.

The DNN receives input features describing the orientation and p_T of each T3, the T5 itself, and the “bridge T3” that connects the three MDs at the center of the T5 (Fig. 6.5a). The coordinates of the “anchor” hit in each MD are also provided—for PS p_T -modules, the hit in the P-layer is taken as the anchor hit, whereas the innermost hit is taken as the anchor hit for 2S modules. There are 38 input features in total: the p_T estimate and radius of a circular fit to the hits in each of the T3s; the (r, ϕ, z) coordinates, η , and layer of each anchor hit; the p_T estimate, η , and ϕ of the T5; and the radius of a circular fit to the hits in the bridge T3. Because the DNN will classify T5s as fake or real, binary cross entropy is used as the loss function.

Given the relatively simple input, and the intense throughput requirements of LST, we selected a simple DNN architecture: a two-layer DNN with 32 nodes per hidden layer. The DNN was trained for over 600 epochs, but the loss plateaued after epoch 400 (Fig. 6.6a). The ROC curve at epoch 500 is shown in Fig. 6.6b; larger architectures were explored, but we found that they did not give significantly better performance. Compared to LST, the DNN is able to reduce

the number of false positives—fake T5s incorrectly classified as real—by nearly 50%.

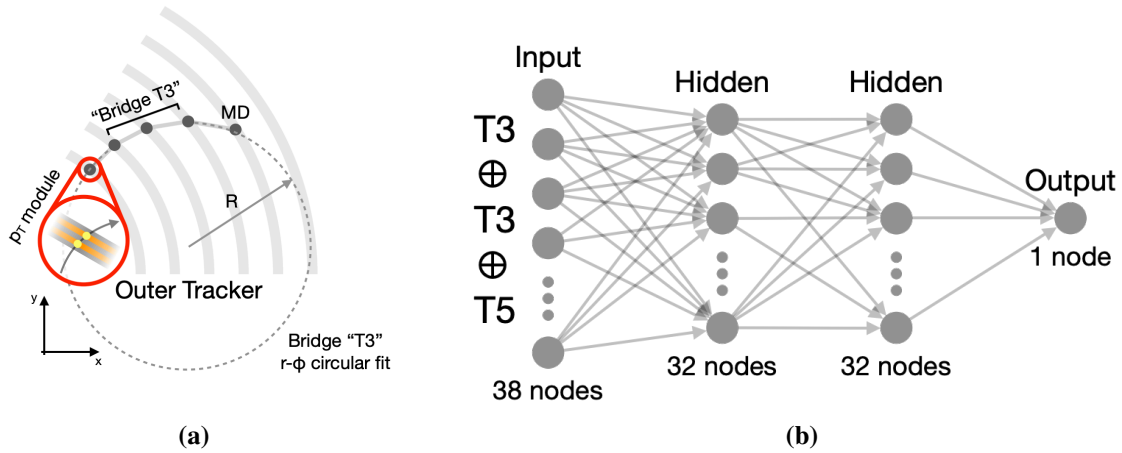


Figure 6.5. The components of a T5 in the r - ϕ plane, including the “bridge” T3, the circular fit, and the anchor hit in an MD (left) alongside the DNN architecture used in this work (right).




Figure 6.6. The average loss after each epoch plotted (a) next to the Receiver Operating Characteristic (ROC) curves for the model after epoch 500 (b). For the ROC curve, the signal efficiency is plotted on the y-axis, while the background efficiency, or fake rate, is plotted on the x-axis.

6.3.2 Results

After verifying that the DNN could classify T5s better than the cuts removed from LST, it was integrated into the LST algorithm so the effect on the final collection of TCs could be evaluated. A working point for the DNN was selected to match the baseline LST true positive rate—the number of T5s correctly classified as real. As such, no loss of efficiency is observed with the DNN integrated into the algorithm (Fig. 6.7). However, the efficiency binned in the displacement of the track tells a different story (Fig. 6.9): the DNN recovers displaced tracks that were being dropped by LST. Meanwhile, the fake rate (Fig. 6.8) is reduced by 40% on average in the barrel, addressing the motivation for this work. These performance enhancements come at no cost in throughput (Fig. 6.10).

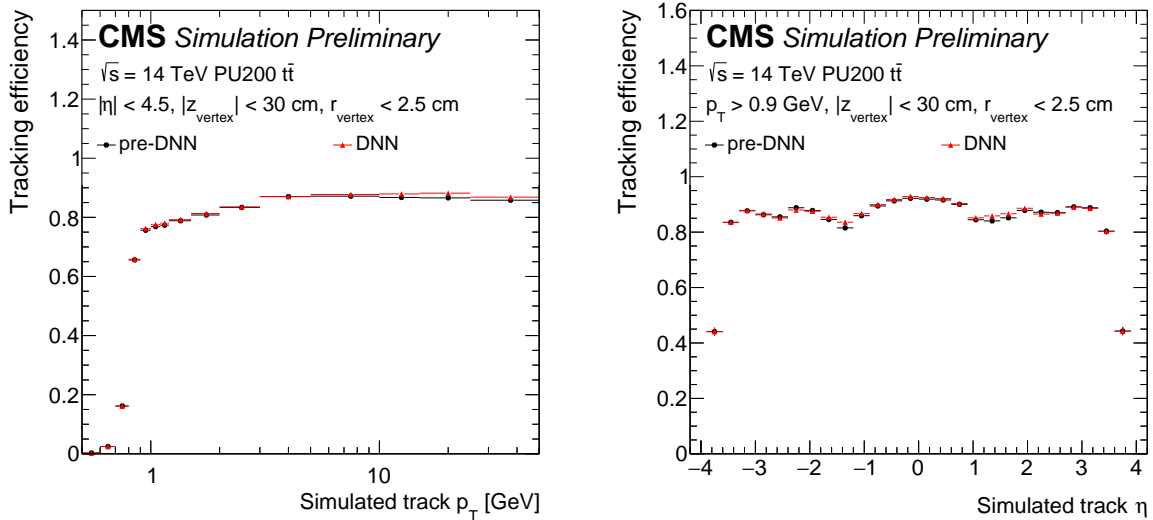


Figure 6.7. The LST efficiency for all TCs plotted as a function of p_T (left) and η (right). The working point for the DNN was selected to match the efficiency of LST, and it is clear that no efficiency is lost.

6.4 Next steps

Immediately, the success of the DNN begs the question, “what is it doing so much better than what was done before?” We are thus currently dissecting the DNN, determining

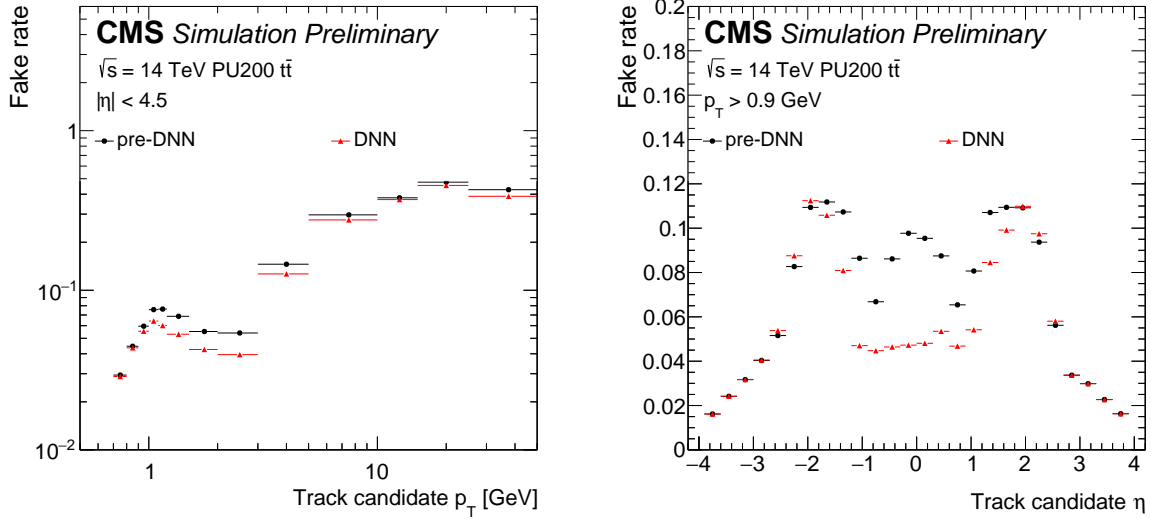


Figure 6.8. The LST fake rate for all TCs plotted as a function of p_T (left) and η (right). Notably, there is a 40% reduction in the fake rate in the barrel, where the T5 fake rate was previously dominant.

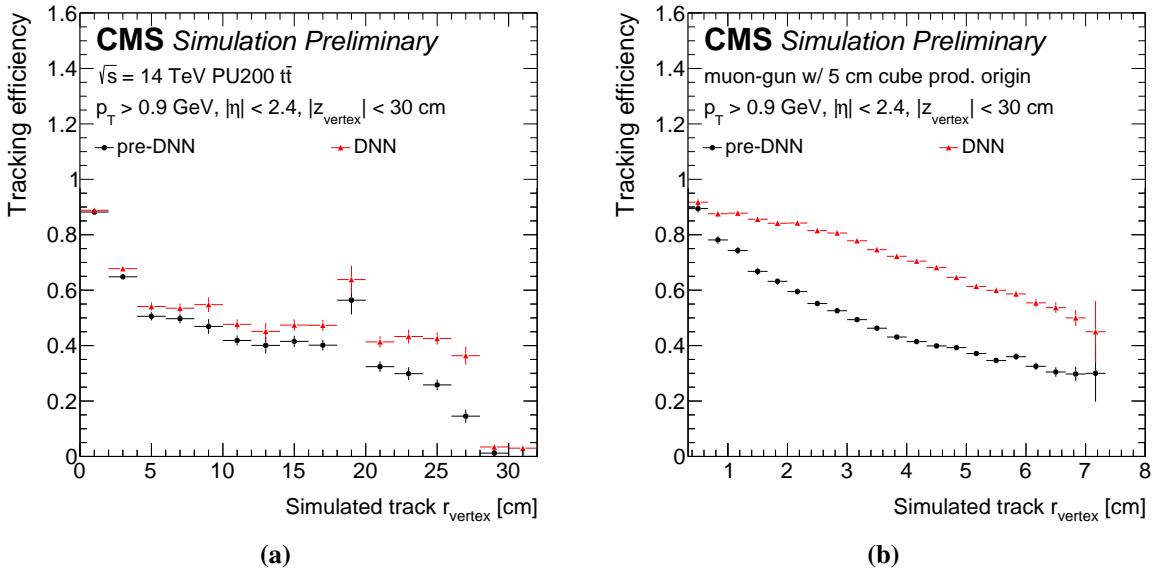


Figure 6.9. The LST efficiency for all TCs plotted as a function of r_{vertex} , i.e. the distance to the production vertex measured in the plane transverse to the beamline. This plot is made with 1000 $t\bar{t}$ events with HL-LHC pile-up (left) and 10,000 "muon-cube" events where muons are produced at points uniformly distributed across a 5 cm cube (right). In both plots, it is clear that the T5-DNN recovers a significant amount of efficiency for displaced tracks.

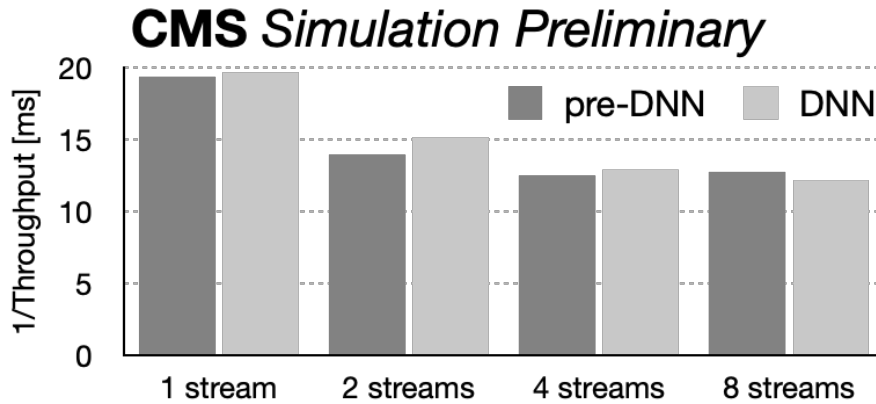


Figure 6.10. The LST throughput plotted for different numbers of parallel CUDA streams that are each opportunistically processing different events concurrently. The throughput is measured before (dark gray) and after (light gray) the DNN is integrated into LST, showing no significant difference.

which input features are most important and understanding how they are used to achieve the observed performance boost. At the same time, there are a number of additional opportunities for leveraging ML to improve LST. The same classification task can be improved for T3s, for instance, and duplicate removal—discerning between overlapping tracks—is a difficult, but vital step where efficiency is critical. Finally, more ambitious ML models could also be envisioned. Recent work suggests that graph neural networks (GNNs) could be used to do end-to-end track-finding and possibly track-fitting as well [143, 144, 145]. We foresee these efforts integrated into LST, where LST provides a fast graph-building step—which is the slowest step in the current GNN tracking efforts—that feeds into a robust GNN-based tracking algorithm.

6.5 Acknowledgements

This chapter is a partial reproduction of the paper “Improving tracking algorithms with machine learning: a case for line-segment tracking at the High Luminosity LHC,” in the proceedings of Connecting the Dots 2023 (arXiv:2403.13166). This work was supported by the U.S. National Science Foundation under Cooperative Agreements OAC-1836650, PHY-2323298, and PHY-2121686 and grant PHY-2209443.

Chapter 7

Exascale Cyberinfrastructure

7.1 The HL-LHC data challenge

In the high luminosity era of the LHC, the CMS Experiment alone is expected to produce roughly 0.5 exabytes (EB) of data every year¹ (Fig. 7.1). The data will need to be distributed across the worldwide LHC computing grid (WLCG), as it is now, to be accessed by thousands of scientists. This defines two important classifications of network traffic: the *distribution* of LHC data, referred to as “production” traffic, and the *access* and final processing of that data, referred to as “user” traffic. Crucially, CMS data movement is entirely driven by the production data transfers, which are centrally managed.

With many PB, soon to be EB, of data that needs to be distributed all over the world, National Research and Education Networks (NRENs) are critical to the continued operation of CMS, and the LHC in general. However, NRENs are entirely opaque to their users—that is, users have no control over the share of the network that they receive—resulting in highly variable networking performance. This kind of unmanaged networking is referred to as “best effort” service, and it already presents significant challenges to CMS operators running production data transfers. For example, high priority data transfers will unpredictably slow down, or even fail, making any precise coordination impossible. These kinds of issues are furthermore difficult

¹For context, the CMS Experiment produced roughly half as much data across the three years of data-taking in Run 2.




Figure 7.1. The HL-LHC tape usage projections at CMS as a function of time, showing data production rates of over 0.5 EB per year after 2030 [139].

to diagnose: with an opaque network, operators must coordinate between the data storage site administrators and network engineers just to gain a preliminary understanding of where the problem may lie. An increase in data volumes by an order of magnitude would inflate these challenges into significant barriers for the data operation at CMS, and the LHC in general. Thus, a novel solution is imperative in order to guarantee the success of the HL-LHC [146, 147].

7.2 Software defined networking

If CMS operators could have a specified bandwidth reserved for a high-priority, production data transfers could be achieved on a more well-defined timescale. Meanwhile, lower priority data transfers, including user traffic, could share the unreserved bandwidth as they do today. This functionality is exactly provided by software defined networking (SDN) [148], which enables the allocation of networking resources in the same way that CPUs and memory are allocated on shared computing clusters. With this model, CMS operators could hold the network accountable for the “promises” (bandwidth guarantees) it has made. By comparing data movement speed

against the network promise and storage site diagnostics, which are already maintained, CMS operators could easily determine, and therefore more quickly address, bottlenecks in throughput. Therefore, by incorporating SDN into the existing data movement infrastructure, CMS would be able to efficiently and reliably move EB of data around the world. Furthermore, while described in the context of CMS, SDN could be used to address the data movement requirements for other LHC experiments like ATLAS, or even other large-scale science experiments.

7.3 Rucio

The distribution of CMS data is managed through Rucio, an open-source data management framework designed specifically for large-scale science [149]. In practice, CMS operators define “rules” which represent the replication of one or more datasets at one or more data storage sites. The operator also assigns a priority to each rule that can be changed at any time. Until the replication is complete, the requisite data transfers are organized for each rule. These data transfers are initialized and managed by the File Transfer Service (FTS) [150], starting with the transfers belonging to the highest priority rule.

7.4 SENSE

The Energy Sciences Network (ESnet), the NREN used by the CMS and ATLAS collaborations in the United States, has developed an SDN product called SENSE: SDN for end-to-end networked science at the exascale [151]. It is designed to provide users with enhanced interactivity via an “intent-based” interface, allowing for the precise management of network resources—the exact functionality needed for HL-LHC data movement. SENSE provides this interactivity through a novel configuration of the data storage sites (Fig. 7.2) that allows data to be accessed across many different IPv6 subnets simultaneously. These subnets logically divide the network traffic, enabling SENSE to allocate bandwidth.

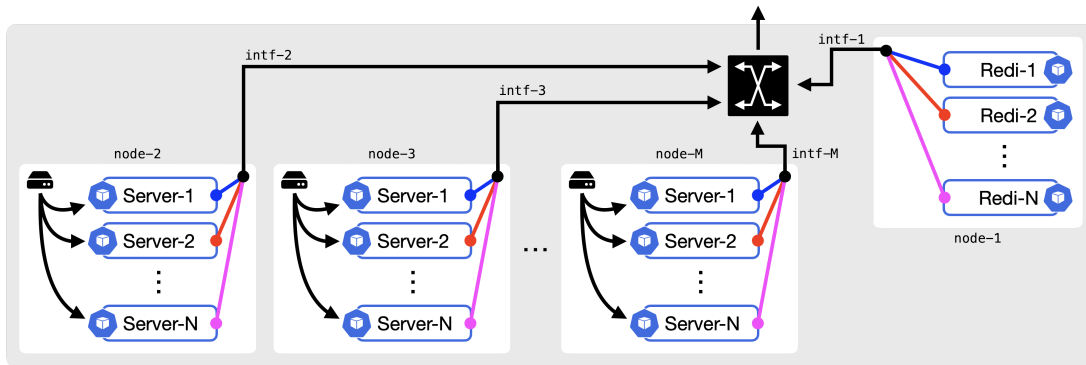


Figure 7.2. A generic site configuration that enables SENSE capabilities, where each “redirector” services a different bandwidth allocation by directing traffic for that allocation to one of the data servers connected to it. Each data server has equal access to the underlying filesystem.

7.5 Rucio-SENSE interoperation model

In order to leverage SENSE for CMS data distribution, the close interoperation of Rucio and SENSE must be implemented with little-to-no changes to Rucio, as it is already used in production. Specifically, priorities assigned to each rule in Rucio should correspond to an appropriate bandwidth reservation made through SENSE, and changes to that rule should result in changes to the reservation. We designed and implemented a Data Movement Manager (DMM) to perform this crucial translation. Importantly, DMM also holds the intelligence for determining what an “appropriate” bandwidth reservation is for a given priority. DMM is therefore the keystone of the Rucio-SENSE interoperation model (Fig. 7.3) first described in Ref. [152]:

1. A Rucio operator initializes a rule with some priority which requests one or more dataset transfers, where each transfer may involve a different pair (source and destination) of sites
2. Rucio sends the following data to DMM for each transfer:
 - Total transfer size
 - Source site
 - Destination site
 - Priority

3. DMM processes the data from Rucio:
 - (a) If the transfer has no priority, place it on best effort service (skip steps below)
 - (b) Reserve an IPv6 address at the source and destination site
 - (c) Compute the bandwidth provision (i.e. promise) appropriate for the transfer priority
4. DMM requests a new promise from SENSE that implements the provisioning from (3c), reprovisioning existing promises where appropriate
5. DMM sends the IPv6 addresses it reserved to Rucio
6. Rucio injects the IPv6 addresses into the FTS request
7. SENSE takes one of the following actions:
 - (a) Begin the construction of a new guaranteed-bandwidth link
 - (b) Do nothing; the transfer will be provided best effort service
8. SENSE sends identifying metadata for the link back to DMM

There are multiple opportunities for optimization in the steps above that we are currently evaluating in our testbed as it evolves. For example, a future implementation of the Rucio-SENSE interoperation model will see the integration of DMM into Rucio, such that steps (2) through (5) can be implemented to better handle a large number of transfers. In addition, the bandwidth provisioning decision, step (3c), is of particular interest. It could be designed to allow for scheduling—e.g. *move Dataset A to Site B in one week*—for instance. Alternatively, transfers could be allocated bandwidth according to their priorities.

7.6 Demonstrations

The Rucio-SENSE interoperation was first demonstrated in a simple test involving two sites, one at UCSD and another at Caltech [153]. These institutions were selected because they




Figure 7.3. A simplified diagram of the Rucio-SENSE interoperation workflow, taken from Ref. [152], with numbered steps: (1) a rule is initialized; (2) Rucio sends transfer description to DMM; (3) DMM translates Rucio request into SENSE provision; (4) DMM sends provision request to SENSE; (5) DMM sends a source IPv6 and destination IPv6 to Rucio; (6) Rucio injects IPv6s from previous step into the FTS request; (7) Either (a) SENSE builds dedicated service or (b) default to best effort service; (8) SENSE sends service metadata to DMM.

each maintain highly active “Tier-2” computing facilities, which are the endpoints of real CMS production data transfers. In order to simulate this kind of network traffic, the IPerf tool was used to emulate activity across best effort network services. At the same time, a private instance of Rucio was deployed, with only UCSD and Caltech as known sites. After the best effort network traffic saturated the 10 Gb/s network link between the two sites, we initialized a rule in Rucio that requested a replica of 750 GB of data stored at UCSD to be made at Caltech. This rule was assigned a high priority. Once the rule was initialized, DMM requested bandwidth reservation of 7 Gb/s for the data transfers. Importantly, the bandwidth received may never fall below a SENSE bandwidth reservation, but it is allowed to exceed it. As shown in Fig. 7.4, the priority data transfers were completed in a timely manner thanks to the large amount of bandwidth reserved for them. The throughput for best effort traffic, meanwhile, was suppressed until the priority transfers finished. This test shows the fundamental action of the Rucio-SENSE interoperation model from end to end, demonstrating that the two technologies could successfully work in concert. Though simple, the test was the result of many months of work, as many aspects of the

implementation of this model are novel, and therefore non-trivial.

A more complicated test was recently performed involving a testbed site at Fermilab (FNAL) in addition to the existing Caltech and UCSD sites. In this test, high-priority data transfers are started between UCSD and Caltech—in particular, they are assigned a priority of 5. Then, data transfers between FNAL and Caltech are initiated with a middling priority of 3. Finally, we began another set of data transfers between UCSD and Caltech with a low priority of 2. All sites had 100 Gb/s connectivity between each other, so the priorities correspond to 50 Gb/s, 30 Gb/s, and 2 Gb/s, respectively—there was no best-effort traffic in this test. As the data moved from source to destination, we swapped the priority assignments every one to two hours. In Fig. 7.5, it can be seen that the bandwidth is appropriately limited for the three data transfers at any one time.




Figure 7.4. Throughput measurements of the first demonstration of the Rucio-SENSE interoperation prototype involving three sites: UCSD and Caltech. Background traffic is simulated with the IPerf tool, emulating activity across best effort network services. Then, a set of priority data transfers are initialized between the two sites with at least 7 Gb/s of requested bandwidth. This throttles the best effort traffic until the priority transfers are complete and the reservation is released.

7.7 Next steps

With the Rucio-SENSE interoperation model now demonstrated successfully multiple times, we are now interested in including more sites in the existing testbed as we walk the project closer to production. In addition, the intelligence of the DMM priority allocation will continue




Figure 7.5. Throughput measurements of a demonstration of the Rucio-SENSE interoperation prototype involving three sites: UCSD, FNAL, and Caltech. Three sets of data transfers are started simultaneously: one from UCSD to Caltech with priority 5, another from FNAL to Caltech with priority 3, and one more from UCSD to Caltech with a priority of 2. Every one to two hours, the priorities are swapped between the three. The throughput for each transfer correctly reflects the appropriate bandwidth for its priority at any one time. In particular, the priorities 5, 3, and 2 correspond to 50 Gb/s, 30 Gb/s, and 20 Gb/s, respectively.

to be refined, with special attention given to the possibility of data transfer scheduling.

7.8 Acknowledgements

This chapter is a partial reproduction of the paper “Managed Network Services for Exascale Data Movement Across Large Global Scientific Collaborations,” in the proceedings of XLOOP 2022 (doi:10.1109/XLOOP56614.2022.00008). This ongoing work is partially supported by the US National Science Foundation (NSF) Grants OAC-2030508, OAC-1841530, OAC-1836650, MPS-1148698, and PHY-1624356. In addition, the development of SENSE is supported by the US Department of Energy (DOE) Grants DE-SC0015527, DE-SC0015528, DE-SC0016585, and FP-00002494. Finally, this work would not be possible without the significant contributions of collaborators at ESnet, Caltech, and SDSC.

Appendix A

CMS coordinate system

The CMS coordinate system (Fig. A.1) takes the z axis to be along the beamline and the origin at the collision point. The x axis points to the center of the LHC and the y axis is orthogonal to the x - z plane. The x - y plane defines the “transverse” plane, wherein quantities like transverse momentum p_T are measured. Naturally, cylindrical coordinates (r, z, ϕ) are used frequently, where r is the transverse distance from the z -axis to any point in space and ϕ is the azimuthal angle. The pseudorapidity η is often used in place of the polar angle θ , where η is defined as

$$\eta \equiv -\ln \left[\tan \left(\frac{\theta}{2} \right) \right] \quad (\text{A.1})$$

That is, η is zero in the transverse plane, and it approaches infinity along either direction of the beamline. It is preferred over θ because differences in η are uniquely Lorentz-invariant under boosts along the z axis, which is important in the context of proton-proton collisions, where the interacting partons will carry some random fraction of the proton’s energy according to the parton distribution function. Moreover, since the pseudorapidity is simply a function of θ , which is easily measured, it is preferred over the actual rapidity, which is a function of the particle’s energy. However, the pseudorapidity of a particle is approximately equal to the rapidity, so long as its rest mass is much smaller than the magnitude of its momentum.




Figure A.1. The CMS coordinate system, from Ref [154], with a cylinder representing the volume of the CMS detector. The origin is taken to be the interaction point (IP), with the z -axis along the beamline, the x -axis towards the center of the LHC, and the y -axis pointing upwards, perpendicular to the x - z plane.

Bibliography

- [1] C. Taylor, *The Atomists: Leucippus and Democritus - Fragments (a text and translation with a commentary)*. Phoenix Supplementary Volumes 36, Phoenix pre-Socratics 5., University of Toronto Press, 1999.
- [2] H. Margenau, *Physics and Philosophy: Selected Essays*. D. Reidel Publishing Company, 1978.
- [3] D. J. Griffiths, *Introduction to elementary particles*. Physics textbook, New York, USA: Wiley, 2nd ed., 2008.
- [4] J. Dalton, “On the proportion of the several gases in the atmosphere,” in *Memoirs of the Literary and Philosophical Society of Manchester*, vol. 1 of 2nd Series, p. 244, 1805.
- [5] J. Dalton, *A New System of Chemical Philosophy*. Strand, London: S. Russell for R. Bickerstaff, 1898.
- [6] M. I. Grossman, “William Higgins at the Dublin Society, 1810–20: the loss of a professorship and a claim to the atomic theory,” *Notes and Records of the Royal Society*, vol. 64, no. 4, pp. 417–434, 2010.
- [7] H. Becquerel, “Sur les radiations émises par phosphorescence,” *Comptes rendus de l’Academie des Sciences, Paris*, vol. 122, p. 420, 1896.
- [8] M. Curie and G. Lippmann M., *Rayons émis par les composés de l’uranium et du thorium*. Gauthier-Villars, 1898.
- [9] P. Radvanyi and J. Villain, “The discovery of radioactivity,” *Comptes Rendus Physique*, vol. 18, no. 9, pp. 544–550, 2017. Science in the making: The Comptes rendus de l’Académie des sciences throughout history.
- [10] J. J. Thompson, “Cathode rays,” *The Philosophical Magazine*, vol. 44, pp. 293–316, 1897.
- [11] J. Perrin, “Brownian movement and molecular reality,” *Nature*, vol. 86, pp. 105–105, Mar 1911. translated from *Annales de Chimie et de Physique*, 8th Series, 1909, by F. Soddy.
- [12] R. Brown, “XXVII. A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the

general existence of active molecules in organic and inorganic bodies,” *The Philosophical Magazine*, vol. 4, no. 21, pp. 161–173, 1828.

- [13] A. Einstein, “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen,” *Annalen der Physik*, vol. 322, no. 8, pp. 549–560, 1905.
- [14] A. Einstein, “Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt,” *Annalen der Physik*, vol. 322, no. 6, pp. 132–148, 1905.
- [15] M. Planck, “On the Theory of the Energy Distribution Law of the Normal Spectrum,” *Verhandlungen der Deutsche Physikalische Gesellschaft Berlin*, vol. 2, p. 237, 1900.
- [16] N. Bohr, “I. On the constitution of atoms and molecules,” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 26, no. 151, pp. 1–25, 1913.
- [17] J. Chadwick, “Possible existence of a neutron,” *Nature*, vol. 129, pp. 312–312, Feb 1932.
- [18] M. Gell-Mann, “The Eightfold Way: A Theory of strong interaction symmetry,” Mar 1961.
- [19] V. E. Barnes, P. L. Connolly, D. J. Crennell, B. B. Culwick, W. C. Delaney, W. B. Fowler, P. E. Hagerty, E. L. Hart, N. Horwitz, P. V. C. Hough, J. E. Jensen, J. K. Kopp, K. W. Lai, J. Leitner, J. L. Lloyd, G. W. London, T. W. Morris, Y. Oren, R. B. Palmer, A. G. Prodell, D. Radojčić, D. C. Rahm, C. R. Richardson, N. P. Samios, J. R. Sanford, R. P. Shutt, J. R. Smith, D. L. Stonehill, R. C. Strand, A. M. Thorndike, M. S. Webster, W. J. Willis, and S. S. Yamamoto, “Observation of a hyperon with strangeness minus three,” *Phys. Rev. Lett.*, vol. 12, pp. 204–206, Feb 1964.
- [20] D. B. Lichtenberg and S. P. Rosen, *Developments in the Quark Theory of Hadrons*. Nonantum, MA: Hadronic Press, 1980. Not held by the CERN library.
- [21] E. D. Bloom, D. H. Coward, H. DeStaebler, J. Drees, G. Miller, L. W. Mo, R. E. Taylor, M. Breidenbach, J. I. Friedman, G. C. Hartmann, and H. W. Kendall, “High-energy inelastic $e - p$ scattering at 6° and 10° ,” *Phys. Rev. Lett.*, vol. 23, pp. 930–934, Oct 1969.
- [22] M. Breidenbach, J. I. Friedman, H. W. Kendall, E. D. Bloom, D. H. Coward, H. DeStaebler, J. Drees, L. W. Mo, and R. E. Taylor, “Observed behavior of highly inelastic electron-proton scattering,” *Phys. Rev. Lett.*, vol. 23, pp. 935–939, Oct 1969.
- [23] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” *Phys. Lett. B*, vol. 716, pp. 1–29, 2012.
- [24] CMS Collaboration, “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC,” *Phys. Lett. B*, vol. 716, pp. 30–61, 2012.

- [25] S. Coleman, *Aspects of Symmetry: Selected Erice Lectures*. Cambridge University Press, 1988.
- [26] J. P. A. Ioannidis, I.-A. Cristea, and K. W. Boyack, “Work honored by nobel prizes clusters heavily in a few scientific fields,” *PLOS ONE*, vol. 15, pp. 1–11, Jul 2020.
- [27] M. G. Holloway and C. P. Baker, “How the barn was born,” *Physics Today*, vol. 25, pp. 9–9, 07 1972.
- [28] M. Herndon. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsCombined>. Accessed: 02-15-2024.
- [29] ATLAS Collaboration, “Standard Model Summary Plots October 2023,” tech. rep., CERN, Geneva, 2023. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2023-039>.
- [30] M. Srednicki, *Quantum Field Theory*. Physics textbook, Cambridge, UK: Cambridge University Press, 1st ed., 2007.
- [31] M. E. Peskin and D. V. Schroeder, *An Introduction To Quantum Field Theory*. Physics textbook, Reading, USA: Addison-Wesley, 1st ed., 1995.
- [32] F. Englert and R. Brout, “Broken symmetry and the mass of gauge vector mesons,” *Phys. Rev. Lett.*, vol. 13, pp. 321–323, Aug 1964.
- [33] P. Higgs, “Broken symmetries, massless particles and gauge fields,” *Physics Letters*, vol. 12, no. 2, pp. 132–133, 1964.
- [34] P. W. Higgs, “Broken symmetries and the masses of gauge bosons,” *Phys. Rev. Lett.*, vol. 13, pp. 508–509, Oct 1964.
- [35] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, “Global conservation laws and massless particles,” *Phys. Rev. Lett.*, vol. 13, pp. 585–587, Nov 1964.
- [36] S. Weinberg, “A Model of Leptons,” *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 1967.
- [37] Y. Nambu and G. Jona-Lasinio, “Dynamical model of elementary particles based on an analogy with superconductivity. II.,” *Phys. Rev.*, vol. 124, pp. 246–254, 1961.
- [38] Particle Physics Project Prioritization Panel, “Pathways to innovation and discovery in particle physics.” <https://www.usparticlephysics.org/2023-p5-report/>, 2023. Accessed: 04-25-2024.
- [39] S. D. Bass, A. De Roeck, and M. Kado, “The higgs boson implications and prospects for future discoveries,” *Nature Reviews Physics*, vol. 3, pp. 608–624, Sep 2021.
- [40] V. C. Rubin, J. Ford, W. K., and N. Thonnard, “Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa→Sc.,” *The Astrophysical Journal*, vol. 225, pp. L107–L111, Nov 1978.

- [41] R. H. Wechsler and J. L. Tinker, “The connection between galaxies and their dark matter halos,” *Annual Review of Astronomy and Astrophysics*, vol. 56, no. Volume 56, 2018, pp. 435–487, 2018.
- [42] Planck Collaboration, “Planck 2018 results - I. Overview and the cosmological legacy of Planck,” *Astronomy & Astrophysics*, vol. 641, p. A1, 2020. Table 7.
- [43] D. B. Kaplan, H. Georgi, and S. Dimopoulos, “Composite Higgs Scalars,” *Phys. Lett. B*, vol. 136, pp. 187–190, 1984.
- [44] N. Arkani-Hamed, A. G. Cohen, E. Katz, and A. E. Nelson, “The Littlest Higgs,” *JHEP*, vol. 07, p. 034, 2002.
- [45] R. Contino, Y. Nomura, and A. Pomarol, “Higgs as a holographic pseudoGoldstone boson,” *Nucl. Phys. B*, vol. 671, pp. 148–174, 2003.
- [46] G. C. Branco, P. M. Ferreira, L. Lavoura, M. N. Rebelo, M. Sher, and J. P. Silva, “Theory and phenomenology of two-Higgs-doublet models,” *Phys. Rept.*, vol. 516, pp. 1–102, 2012.
- [47] J. F. Gunion and H. E. Haber, “Conditions for CP-violation in the general two-Higgs-doublet model,” *Phys. Rev. D*, vol. 72, p. 095002, 2005.
- [48] B. Grzadkowski, O. M. Ogreid, and P. Osland, “Natural Multi-Higgs Model with Dark Matter and CP Violation,” *Phys. Rev. D*, vol. 80, p. 055013, 2009.
- [49] L. Susskind, “Dynamics of Spontaneous Symmetry Breaking in the Weinberg-Salam Theory,” *Phys. Rev. D*, vol. 20, pp. 2619–2625, 1979.
- [50] G. 't Hooft, “Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking,” *NATO Sci. Ser. B*, vol. 59, pp. 135–157, 1980.
- [51] Service graphique, CERN, “Overall view of the LHC.” <https://cds.cern.ch/record/1708849>, 2014. General Photo.
- [52] M. Brice, “Aerial View of the CERN taken in 2008.” <https://cds.cern.ch/record/1295244>, 2008. General Photo.
- [53] E. Lopienska, “The CERN accelerator complex, layout in 2022.” <https://cds.cern.ch/record/2800984>, 2022. General Photo.
- [54] D. Dominguez, “3D cut of the LHC dipole.” <https://cds.cern.ch/record/1741036>, 2014. General Photo.
- [55] L. Taylor, “CMS Document 5592-v1.” <https://cms-docdb.cern.ch/cgi-bin/PublicDocDB/ShowDocument?docid=5592>. Accessed: 04-21-2024.
- [56] CERN, “Interim summary report on the analysis of the 19 September 2008 incident at the LHC,” tech. rep., CERN, Geneva, 2008.

- [57] W. Herr and B. Muratori, “Concept of luminosity,” 2006.
- [58] CMS Collaboration, “Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 13$ TeV,” tech. rep., CERN, Geneva, 2016.
- [59] CMS Collaboration, “Public CMS Luminosity Information.” <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>, 2024. Accessed: 04-16-2024.
- [60] I. B. Alonso, O. Brüning, P. Fessia, M. Lamont, L. Rossi, L. Tavian, and M. Zerlauth, *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs, Geneva: CERN, 2020.
- [61] CERN, “The Large Hadron Collider.” <https://lhcb.cern>. ATLAS webpage. Accessed: 04-21-2024.
- [62] ATLAS Collaboration, “ATLAS.” <https://atlas.cern/detector>. ATLAS webpage. Accessed: 04-21-2024.
- [63] CMS Collaboration, “CMS.” <https://cms.cern/detector>. CMS webpage. Accessed: 04-21-2024.
- [64] T. Sakuma, “Cutaway diagrams of CMS detector.” <https://cds.cern.ch/record/2665537>, 2019.
- [65] M. Brice, “Replacement of the heart of the CMS experiment - the pixel detector, Part2.” <https://cds.cern.ch/record/2254952>, 2017. General Photo.
- [66] M. Brice, M. Hoch, and J. Gobin, “View of the CMS Detector before closure.” <https://cds.cern.ch/record/1133594>, 2008.
- [67] CMS Collaboration, “Commissioning of the CMS Experiment and the Cosmic Run at Four Tesla,” *JINST*, vol. 5, p. T03001, 2010.
- [68] CMS Collaboration, “Bending particles.” <https://cms.cern/detector/bending-particles>, 2024. Accessed: 04-25-2024.
- [69] Anonymous, “Solenoid and ampère’s law.” https://commons.wikimedia.org/wiki/File:Solenoid_and_Ampere_Law_-_2.png, 2018. Accessed: 05-28-2024.
- [70] CMS Collaboration, “Precise Mapping of the Magnetic Field in the CMS Barrel Yoke using Cosmic Rays,” *JINST*, vol. 5, p. T03021, 2010.
- [71] CMS Collaboration, *The CMS magnet project: Technical Design Report*. Technical design report. CMS, Geneva: CERN, 1997.
- [72] Physical Therapy, “The Fatal Current,” *Physical Therapy*, vol. 46, pp. 968–969, Sept 1966.

- [73] A. Berger, “How does it work?: Magnetic resonance imaging,” *BMJ*, vol. 324, pp. 35–35, Jan 2002.
- [74] D. Campi and D. Stickland, “CMS: a super solenoid is ready for business,” *CERN Courier*, vol. 47, p. P10003, Mar 2017.
- [75] OPERA Collaboration, “The OPERA experiment in the CERN to Gran Sasso neutrino beam,” *JINST*, vol. 4, p. P04018, 2009.
- [76] CDF Collaboration, “CDF central outer tracker,” *Nucl. Instrum. Meth. A*, vol. 526, pp. 249–299, 2004.
- [77] CMS Collaboration, “Tracking.” <https://cms.cern/detector/identifying-tracks>, 2024. Accessed: 04-25-2024.
- [78] M. Hoch, “Work on CMS Tracker during Long Shutdown 1 (LS1) - CMS Forward Pixel Installation.” <https://cds.cern.ch/record/2017264>, 2014. CMS Collection.
- [79] M. Brice, “First half of CMS inner tracker barrel.” <https://cds.cern.ch/record/995912>, 2006.
- [80] T. Sakuma, “SketchUp Drawings of the Phase 1 pixel detector alongside the current pixel detector,” 2012.
- [81] CMS Collaboration, “Energy of electrons and photons (ECAL).” <https://cms.cern/index.php/detector/measuring-energy/energy-electrons-and-photons-ecal>, 2024. Accessed: 04-27-2024.
- [82] CMS Collaboration, *The CMS electromagnetic calorimeter project: Technical Design Report*. Technical design report. CMS, Geneva: CERN, 1997.
- [83] M. Brice, “Images of the CMS ECAL Barrel.” <https://cds.cern.ch/record/1431477>, 2008. CMS Collection.
- [84] CMS Collaboration, *The CMS hadron calorimeter project: Technical Design Report*. Technical design report. CMS, Geneva: CERN, 1997.
- [85] H. Kolanoski and N. Wermes, *Particle Detectors*. Physics textbook, Oxford, UK: Oxford University Press, 1st ed., 2020.
- [86] CMS Collaboration, “Energy of hadrons (HCAL).” <https://cms.cern/detector/measuring-energy/energy-hadrons-hcal>, 2024. Accessed: 04-27-2024.
- [87] M. Brice, L. Vaillat, and L. Lazic, “Images of the CMS HCAL Barrel.” <https://cds.cern.ch/record/1431485>, 2008. CMS Collection.
- [88] P. Ginter, “Workers in Mormansk sitting on decommissioned WWII artillery shell casings from the Russian Northern Fleet.” <https://www.peterginter.com/pg-portfolio/cern-lhc/>, 2014. Accessed: 04-22-2024.

- [89] CMS Collaboration, *The CMS muon project: Technical Design Report*. Technical design report. CMS, Geneva: CERN, 1997.
- [90] CMS Collaboration, “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV,” *JINST*, vol. 13, no. 06, p. P06015, 2018.
- [91] CMS Collaboration, “Detecting muons.” <https://cms.cern/detector/detecting-muons>, 2024. Accessed: 04-27-2024.
- [92] M. Hoch, “Early Construction: Muon Chambers.” <https://cds.cern.ch/record/1274451>, 2007. CMS Collection.
- [93] CMS Collaboration, “Muon drift tubes.” <https://cms.cern/index.php/detector/detecting-muons/muon-drift-tubes>, 2024. Accessed: 04-27-2024.
- [94] CERN, “The High-Luminosity LHC (HL-LHC).” <https://voisins.web.cern.ch/en/high-luminosity-lhc-hl-lhc>, 2024. Accessed: 04-02-2024.
- [95] CMS Collaboration and T. McCauley, “Collision events recorded by CMS in 2016.” <https://cds.cern.ch/record/224114>. CMS Collection, CMS-PHO-EVENTS-2017-001.
- [96] CMS Collaboration and T. McCauley, “Collisions recorded by the CMS detector on 14 Oct 2016 during the high pile-up fill.” <https://cds.cern.ch/record/2231915>. CMS Collection, CMS-PHO-EVENTS-2016-008.
- [97] CMS Collaboration, “The Phase-2 Upgrade of the CMS Tracker,” tech. rep., CERN, Geneva, 2017.
- [98] CMS Collaboration, “Technical proposal for a MIP timing detector in the CMS experiment Phase 2 upgrade,” tech. rep., CERN, Geneva, 2017.
- [99] CMS Collaboration, “The Phase-2 Upgrade of the CMS Endcap Calorimeter,” tech. rep., CERN, Geneva, 2017.
- [100] CMS Collaboration, “The Phase-2 Upgrade of the CMS Barrel Calorimeters,” tech. rep., CERN, Geneva, 2017. This is the final version, approved by the LHCC.
- [101] CERN, “Copyright and terms of use of CERN content.” <https://copyright.web.cern.ch/>, 2024. Accessed: 04-25-2024.
- [102] CMS Collaboration, “A portrait of the Higgs boson by the CMS experiment ten years after the discovery,” *Nature*, vol. 607, no. 7917, p. 60, 2022.
- [103] ATLAS Collaboration, “A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery,” *Nature*, vol. 607, pp. 52–59, Jul 2022.
- [104] LHC Higgs Cross Section Working Group, “Higgs Properties: Report of the LHC Higgs Cross Section Working Group,” in *Handbook of LHC Higgs Cross Sections* (S. Heinemeyer, C. Mariotti, G. Passarino, and R. Tanaka, eds.), no. 3 in CERN Yellow Reports: Monographs, CERN, Jul 2013.

- [105] LHC Higgs Cross Section Working Group, “Deciphering the nature of the Higgs sector,” in *Handbook of LHC Higgs Cross Sections* (D. de Florian, C. Grojean, F. Maltoni, C. Mariotti, A. Nikitenko, M. Pieri, P. Savard, M. Schumacher, and R. Tanaka, eds.), vol. 2 of *CERN Yellow Reports: Monographs*, CERN, May 2017.
- [106] B. Grzadkowski, M. Iskrzyński, M. Misiak, and J. Rosiek, “Dimension-six terms in the standard model lagrangian,” *Journal of High Energy Physics*, vol. 2010, p. 85, Oct 2010.
- [107] M. Rauch, “Vector-boson fusion and vector-boson scattering,” 2016.
- [108] B. Henning, D. Lombardo, M. Riembau, and F. Riva, “Measuring Higgs couplings without Higgs bosons,” *Phys. Rev. Lett.*, vol. 123, p. 181801, Oct 2019.
- [109] C. Polly, “First results from the Muon $g-2$ Experiment at Fermilab.” https://www.fermilabcommunity.org/wp-content/uploads/2021/05/Polly_210527_CAB_final.pdf, 2021. Slide 17.
- [110] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector,” *JINST*, vol. 12, p. P10003, 2017.
- [111] S. R. Davis, “Interactive Slice of the CMS detector.” <https://cds.cern.ch/record/2205172>, 2016.
- [112] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_T jet clustering algorithm,” *JHEP*, vol. 04, p. 063, 2008.
- [113] M. Cacciari, G. P. Salam, and G. Soyez, “FastJet user manual,” *Eur. Phys. J. C*, vol. 72, p. 1896, 2012.
- [114] E. Bols, J. Kieseler, M. Verzetti, M. Stoye, and A. Stakia, “Jet Flavour Classification Using DeepJet,” *JINST*, vol. 15, no. 12, p. P12012, 2020.
- [115] N. Bartosik, “B-tagging diagram.” https://commons.wikimedia.org/wiki/File:B-tagging_diagram.png, 2024. Accessed: 04-19-2024.
- [116] H. Qu and L. Gouskos, “ParticleNet: Jet Tagging via Particle Clouds,” *Phys. Rev. D*, vol. 101, no. 5, p. 056019, 2020.
- [117] CMS Collaboration, “Displays of candidate events in the search for new heavy resonances decaying to dibosons in the all-jets final state in the CMS detector.” <https://cds.cern.ch/record/2809446>, 2022. CMS Collection.
- [118] D. Sandlin, “Bullets HITTING Bullets in Slow Motion - THE IMPOSSIBLE SHOT - Smarter Every Day 287.” <https://youtu.be/tcQVrD7RnNI?si=IRUzaQTSJNhL8MW8>, 2023.

- [119] A. Buckley, J. Butterworth, S. Gieseke, D. Grellscheid, S. Höche, H. Hoeth, F. Krauss, L. Lönnblad, E. Nurse, P. Richardson, S. Schumann, M. H. Seymour, T. Sjöstrand, P. Skands, and B. Webber, “General-purpose event generators for LHC physics,” *Phys. Rept.*, vol. 504, pp. 145–233, 2011.
- [120] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert, and J. Winter, “Event generation with SHERPA 1.1,” *JHEP*, vol. 02, p. 007, 2009.
- [121] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations,” *JHEP*, vol. 07, p. 079, 2014.
- [122] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, “An introduction to PYTHIA 8.2,” *Comput. Phys. Commun.*, vol. 191, p. 159, 2015.
- [123] S. Alioli, P. Nason, C. Oleari, and E. Re, “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX,” *JHEP*, vol. 06, p. 043, 2010.
- [124] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with parton shower simulations: the POWHEG method,” *JHEP*, vol. 11, p. 070, 2007.
- [125] S. Agostinelli *et al.*, “GEANT4—a simulation toolkit,” *Nucl. Instrum. Meth. A*, vol. 506, p. 250, 2003.
- [126] CMS Collaboration, “CMS Monte Carlo production overview.” <https://opendata.cern.ch/docs/cms-mc-production-overview>, 2024. Accessed: 04-16-2024.
- [127] WLCG, “Worldwide LHC Computing Grid.” <https://wlcg.web.cern.ch>. WLCG webpage. Accessed: 04-21-2024.
- [128] C. Collaboration, “The CMS statistical analysis and combination tool: COMBINE,” 2024.
- [129] C. H. de Lima, D. Stolarski, and Y. Wu, “Status of negative coupling modifiers for extended Higgs sectors,” *Phys. Rev. D*, vol. 105, p. 035019, Feb 2022.
- [130] I. Low and J. Lykken, “Revealing the electroweak properties of a new scalar resonance,” *Journal of High Energy Physics*, vol. 2010, p. 53, Oct 2010.
- [131] H. Georgi and M. Machacek, “Doubly charged higgs bosons,” *Nuclear Physics B*, vol. 262, no. 3, p. 463, 1985.
- [132] CMS Collaboration, “Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13$ TeV,” *The European Physical Journal C*, vol. 79, p. 421, May 2019.

- [133] D. Stolarski and Y. Wu, “Tree-level interference in vector boson fusion production of Vh ,” *Phys. Rev. D*, vol. 102, no. 3, p. 033006, 2020.
- [134] CMS Collaboration, “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector,” *Journal of Instrumentation*, vol. 14, p. P07004, Jul 2019.
- [135] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics,” *Eur. Phys. J. C*, vol. 71, p. 1554, 2011. [Erratum: 10.1140/epjc/s10052-013-2501-z].
- [136] ATLAS and CMS Collaborations, “Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV,” *JHEP*, vol. 08, p. 045, 2016.
- [137] G. Kasieczka, B. Nachman, M. D. Schwartz, and D. Shih, “Automating the abcd method with machine learning,” *Phys. Rev. D*, vol. 103, p. 035021, Feb 2021.
- [138] CMS Collaboration, “Search for pair-produced vector-like leptons in final states with third-generation leptons and at least three b quark jets in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Phys. Lett. B*, vol. 846, p. 137713, 2023.
- [139] CMS Offline Software and Computing, “CMS Phase-2 Computing Model: Update Document,” tech. rep., CERN, Geneva, 2022.
- [140] G. Cerati, P. Elmer, P. Gattung, L. Giannini, M. Kortelainen, V. Krutelyov, S. Lantz, M. Masciovecchio, T. Reid, A. R. Hall, D. Riley, M. Tadel, E. Vourliotis, P. Wittich, and A. Yagil, “Generalizing mkFit and its application to HL-LHC,” 2023.
- [141] CMS Collaboration, “Search for inelastic dark matter in events with two displaced muons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Phys. Rev. Lett.*, vol. 132, no. 4, p. 041802, 2024.
- [142] CMS Collaboration, “Improved performance of line segment tracking using machine learning.” <https://cds.cern.ch/record/2872904>, 2023.
- [143] X. Ju, D. Murnane, P. Calafiura, N. Choma, S. Conlon, S. Farrell, Y. Xu, M. Spiropulu, J.-R. Vlimant, A. Aurisano, J. Hewes, G. Cerati, L. Gray, T. Klijsma, J. Kowalkowski, M. Atkinson, M. Neubauer, G. DeZoort, S. Thais, A. Chauhan, A. Schuy, S.-C. Hsu, A. Ballou, and A. Lazar, “Performance of a geometric deep learning pipeline for HL-LHC particle tracking,” *The European Physical Journal C*, vol. 81, p. 876, Oct 2021.
- [144] G. DeZoort, P. W. Battaglia, C. Biscarat, and J.-R. Vlimant, “Graph neural networks at the Large Hadron Collider,” *Nature Reviews Physics*, vol. 5, pp. 281–303, May 2023.
- [145] K. Lieret and G. DeZoort, “An object condensation pipeline for charged particle tracking at the high luminosity lh,” 2023.

- [146] HEP Software Foundation, “A Roadmap for HEP Software and Computing R&D for the 2020s,” *Comput. Softw. Big Sci.*, vol. 3, no. 1, p. 7, 2019.
- [147] J. Zurawski, B. Brown, D. Carder, E. Colby, E. Dart, K. Miller, A. Patwa, K. Robinson, L. Rotman, and A. Wiedlea, “2020 High energy physics network requirements review final report,” Tech. Rep. LBNL-2001398, Lawrence Berkeley National Laboratory, Jun 2021.
- [148] K. Benzekki, A. El Fergougui, and A. Elbelrhiti Elalaoui, “Software-defined networking (sdn): a survey,” *Security and Communication Networks*, vol. 9, no. 18, pp. 5803–5833, 2016.
- [149] M. Barisits, T. Beermann, F. Berghaus, B. Bockelman, J. Bogado, D. Cameron, D. Christidis, D. Ciangottini, G. Dimitrov, M. Elsing, V. Garonne, A. di Girolamo, L. Goossens, W. Guan, J. Guenther, T. Javurek, D. Kuhn, M. Lassnig, F. Lopez, N. Magini, A. Molfetas, A. Nairz, F. Ould-Saada, S. Prenner, C. Serfon, G. Stewart, E. Vaandering, P. Vasileva, R. Vigne, and T. Wegner, “Rucio: Scientific data management,” *Computing and Software for Big Science*, vol. 3, p. 11, Aug 2019.
- [150] A. A. Ayllon, M. Salichos, M. K. Simon, and O. Keeble, “FTS3: New Data Movement Service For WLCG,” *J. Phys. Conf. Ser.*, vol. 513, p. 032081, 2014.
- [151] I. Monga, C. Guok, J. MacAuley, A. Sim, H. Newman, J. Balcas, P. DeMar, L. Winkler, T. Lehman, and X. Yang, “Software-defined network for end-to-end networked science at the exascale,” *Future Generation Computer Systems*, vol. 110, pp. 181–201, 2020.
- [152] J. Guiang, A. Arora, D. Davila, J. Graham, D. Mishin, I. Sfiligoi, F. Würthwein, T. Lehman, X. Yang, C. Guok, H. Newman, J. Balcas, and T. Hutton, “Integrating End-to-End Exascale SDN into the LHC Data Distribution Cyberinfrastructure,” in *Practice and Experience in Advanced Research Computing*, 2022.
- [153] F. Würthwein, J. Guiang, A. Arora, D. Davila, J. Graham, D. Mishin, T. Hutton, I. Sfiligoi, H. Newman, J. Balcas, T. Lehman, X. Yang, and C. Guok, “Managed network services for exascale data movement across large global scientific collaborations,” in *4th Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing*, Sep 2022.
- [154] I. Neutelings, “CMS coordinate system.” https://tikz.net/axis3d_cms/. Accessed: 05-30-2024.